To Have or Not to Have: Understanding Wealth Inequality*

Pavel Brendler†
University of Nottingham

Moritz Kuhn
University of Mannheim,
CEPR, and IZA

Ulrike Steins
Evonauts GmbH

September 2, 2025

Abstract

Differences in household savings behavior across income levels are key for understanding wealth inequality. We propose asset market access as a novel explanation for these differences. We document three new facts about the joint distribution of income and wealth in the United States: (i) wealth gaps by income are primarily driven by differences in asset market participation rather than invested amounts; (ii) participation heterogeneity mirrors differences in observed market access; and (iii) most wealth is accumulated through financial products with contractual savings flows, such as mortgage payments and retirement contributions. We then develop a life-cycle model of financial lifetime decisions: buying a home, starting a retirement plan, and becoming an entrepreneur. In the model, income-dependent access to mortgages and retirement plans endogenously determines household participation decisions, while financial contracts and capital gains govern wealth accumulation among the market participants. Equalizing initially heterogeneous access increases wealth for the bottom half of the income distribution by 32%. We evaluate a currently proposed policy to expand access to retirement plans through a government-sponsored option and find it effective in fostering broad-based wealth accumulation. Thus, we conclude that today's asset markets in the United States do not offer a level playing field for wealth building.

Keywords: Wealth Inequality; Labor Market Heterogeneity; Household Portfolios.
JEL codes: D31; E21; H31.

*We are thankful for very useful comments to Árpád Ábrahám, Marco Bassetto, Corina Boar, Spiros Bougheas, Jake Bradley, Anton Braun, Chris Busch (discussant), Jan Eeckhout, Giulio Fella, Teresa Ghilarducci, Benjamin Glasner, Jonathan Heathcote, Joachim Hubmer, Winfried Koeniger, Markus Kondziella (discussant), Dirk Krueger, Felix Kubler, Burhan Kuruscu, Benjamin Moll, Simon Mongey, Ed Nosal, Krzysztof Pytka, Siavash Radpour, José-Víctor Ríos-Rull, Florian Scheuer, Matthias Schief, Minchul Yum, and the seminar participants at the 2024 Arne Ryde Workshop in Lund, 2024 CEPR European Summer Symposium in International Macroeconomics in Paris, 2023 Macro Workshop in Lisbon, 2023 Workshop on "Macroeconomic implications of housing, household finances, and wealth dynamics" in St. Gallen, 2023 SED Meeting in Cartagena, 2022 Econometric Society Meeting in Berlin, Bank of Portugal, Danmarks Nationalbank, Federal Reserve Banks of Atlanta, St. Louis, and Cleveland, Toulouse School of Economics, London School of Economics, Bilkent University, and the Universities of Birmingham, Bologna, Bonn, Dortmund, Frankfurt, Lugano, Nottingham, and Zurich. We thank Moritz Scheidenberger for excellent research assistance. Steins gratefully acknowledges financial support from the Science Foundation of Sparkassen-Finanzgruppe. Kuhn gratefully acknowledges financial support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1–390838866.

†Corresponding author. Address: University of Nottingham, School of Economics, University Park, Nottingham, NG7 2RD. Email: pavel.brendler at nottingham.ac.uk.

1 Introduction

Household saving behavior is a key driver of wealth inequality. Understanding the reasons behind variation in savings and wealth accumulation across households is, therefore, a long-standing topic of economic research. This paper offers a novel perspective on why some households save more than others by combining empirical and theoretical analyses. While existing theories focus on households' willingness to save, we propose differences in the opportunities to save as an alternative explanation. We argue that unequal access to wealth-building opportunities between households with different incomes is an important, yet unexplored, cause of wealth inequality. These access differences are evident in key lifetime savings choices such as buying a home or starting a retirement plan, as some households have their mortgage applications rejected or lack access to employer-sponsored retirement plans. By shifting the focus from willingness to save to market access, we address a fundamental question: Do today's asset markets offer a level playing field for building wealth?

We address this question in the context of the United States, focusing on home equity, employer-sponsored pension plans, and business wealth as primary components of household wealth. Our empirical analysis examines the joint distribution of these asset classes with income and yields three main findings. First, participation in these asset markets rises sharply with income. This finding generalizes the established results of Guiso and Sodini (2013), who document differences in asset market participation by wealth and emphasize their importance for observed wealth differences. Second, we provide evidence that this heterogeneity in participation by income is strongly linked to the differences in asset market access, specifically, mortgage approval rates and offer rates for employer-sponsored retirement plans along the income distribution. This finding supports the conjecture in Dynan et al. (2004) that unequal market access is a key contributor to differences in savings behavior and wealth accumulation. Finally, we document that the savings flows generated by the long-term financial contracts for mortgages and retirement plans account for 70% of household savings in the US economy, underscoring their central role in understanding the saving behavior and wealth inequality.

Guided by these empirical facts, we develop a theory of life-cycle wealth accumulation centered around the key lifetime decisions that impact wealth building: buying a home, starting a retirement plan, and becoming an entrepreneur, as well as the associated financial obligations (mortgage payments and retirement contributions). Our model features heterogeneous asset market access in the form of income-dependent mortgage approval and offer rates for employer-sponsored retirement plans that endogenously determines asset market participation. Capital gains from rising asset prices amplify the wealth accumulation of asset market participants. We demonstrate that this model matches the novel facts on the joint distribution of income and wealth. Using the model, we assess the role of the empirically documented heterogeneity in asset market access for wealth accumulation. We find that heterogeneous asset market access and the resulting differences in asset market participation are key drivers of low savings among income-poor households. Eliminating heterogeneity in access increases wealth accumulation in the bottom half of the income distribution by 32% and decreases it only slightly (3%) in the top half, thus reducing overall wealth inequality.

Unequal access to asset markets is also an important issue in the current policy debate in the United States. Using our model, we evaluate the recently proposed bipartisan *Retirement Saving for Americans Act*, which aims to improve retirement savings access for low-income workers by offering a government-sponsored plan. We find that implementing such a policy successfully helps low-income households accumulate more wealth. Thus, our empirical and theoretical results suggest that asset markets in the United States today do not offer an even playing field for building wealth, providing worse opportunities to low-income households.

The main data source for our empirical analysis is the *Survey of Consumer Finances* (SCF), covering the years 2010 to 2019. The three major portfolio components of US households, home equity, retirement accounts, and business equity, account for 70% of US households' marketable wealth. We refer to their sum as *core wealth*.[1] To characterize the joint distribution of wealth and income, we group households into income deciles and construct wealth-to-income ratios for each group. We then study the variation of wealth-to-income ratios along the income distribution. By averaging wealth within income groups, our analysis abstracts from variation in wealth *within* these groups, but we also demonstrate that between-group inequality accounts for the larger part of total wealth inequality in the data. We establish three new facts about the joint distribution of income and wealth in the United States.

First, we document that wealth-to-income ratios increase sharply with household income for each of the core wealth components. The majority of these differences in wealth by income are driven by differences in asset market participation (extensive margin). Strikingly, we find that accumulated amounts conditional on participation (intensive margin) remain constant across the income distribution, and in the case of home equity, actually decrease with income. These results demonstrate that understanding wealth inequality is a matter of *to have or not to have*.

Second, we examine the sources of the documented heterogeneity in asset market participation and find large differences in observed asset market access by income. Approval of a mortgage application by a bank is the key condition for access to the housing market, and we document substantial heterogeneity in approval rates along the income distribution, increasing by a factor of five between the bottom and top income deciles. Quantitatively, this pattern is consistent with the documented rising participation rate for home equity, suggesting that approval rates, and hence market access, are an important driver of participation rates. In contrast, the share of mortgage applications remains virtually constant along the income distribution, suggesting that differences in demand are unlikely to account for rising wealth-to-income ratios.

Similarly, we document large heterogeneity in access to employer-sponsored retirement accounts. We find employer offer rates to increase by a factor of four along the income distribution. This increase in access is twice as large as the variation in the pick-up rates of these plans. Hence, while the differences in pick-up rates contribute in part to the variation in participation rates, the main source of variation comes from offer rates and, hence, asset market access, as in the case of housing. We provide extensive sensitivity analyses for these facts. The relationship between income and mortgage approval rates or offer rates remains statistically and economically significant even after controlling for a wide range of household characteris-

[1] Core wealth overlaps to a large extent with "illiquid assets" in Kaplan and Violante (2014).

tics. We confirm this finding in other data sources such as the administrative *Home Mortgage Disclosure Act* (HMDA) data on mortgage approval rates and the *National Compensation Survey* provided by the Bureau of Labor Statistics for pension plans.

In the final step of our empirical analysis, we highlight the importance of contractual savings flows for household wealth accumulation. Specifically, we quantify the proportion of total household savings in the US economy accounted for by mortgage payments (net of interest payments) and contributions to employer-sponsored retirement accounts (including employer contributions). Using aggregate data from the National Income and Product Accounts and SCF microdata, we estimate that contractual savings flows associated with mortgages and pension plans account for 70% of the aggregate savings rate. This finding points to a key role of long-term financial contracts for understanding saving behavior and wealth accumulation.

Motivated by these empirical findings, we develop a model featuring key financial lifetime decisions: buying a house, starting a retirement plan, and becoming an entrepreneur. Our framework differs from existing models of life-cycle wealth accumulation in two important ways. First, as opposed to many macroeconomic models, in which wealth accumulation is governed by the Euler equation and marginal investment decisions, our model focuses on the discrete choices for these lifetime financial decisions.[2] Conditional on the decision to save, participants in the housing and pension markets commit to the fulfillment of predetermined financial contracts (e.g., fixed mortgage contracts and pension plans), behaving like wealthy hand-to-mouth agents in Kaplan and Violante (2014) and Bayer et al. (2019). We abstract from saving in liquid assets and short-run consumption smoothing because we document that non-core wealth contributes little to life-cycle wealth accumulation. As a second departure from the existing literature, we introduce heterogeneity in access to asset markets. Following Luetticke (2021) and Bilbiie (2024), we model trading opportunities in frictional financial markets as stochastic. Guided by our empirical evidence, we allow these stochastic trading opportunities to differ with income. Depending on their income, households face different probabilities of being approved for a mortgage or being offered a pension plan. Thus, we follow the incomplete markets tradition (Bewley, 1986; Huggett, 1993; Aiyagari, 1994) and abstract from microfounding the lack of asset market access but focus on its consequences.[3]

We calibrate the model to the US economy. Our first important finding is that the model with its focus on lifetime financial decisions and stylized wealth accumulation contracts closely matches the joint distribution of income and wealth. Although the calibration does not target heterogeneity along the extensive or intensive margin of any portfolio component, it still achieves an excellent fit for both margins along the income distribution. Consequently, it also closely matches the empirically documented rising wealth-to-income ratios. While the model closely matches between-income-group wealth inequality, it understates within-income-group wealth inequality because it abstracts from other mechanisms of wealth heterogeneity, such as precautionary savings or idiosyncratic return differences. Finally, the model also closely matches the

[2]To account for the important role of business wealth as part of core wealth, we introduce a discrete occupational choice for entrepreneurial activity.

[3]For entrepreneurial activity, we focus on the friction of finding a valuable business idea at the occupational decision stage rather than on the financial constraint of scaling up a valuable idea. In the spirit of the housing and retirement markets, market entry (receiving an idea) is the key friction for entrepreneurs (Bloom et al., 2020).

untargeted increasing savings rates by income (Dynan et al., 2004). The close model fit offers a new interpretation of these differences as arising from differences in market access by income, providing evidence for the validity of this conjecture in Dynan et al. (2004, p. 436).

To quantify the importance of income-dependent access to asset markets for wealth accumulation and inequality, we conduct an experiment in which we remove income dependence in asset market access. We do this by setting mortgage approval rates and employer offer rates for pension plans to their respective economy-wide averages. We find that participation in the housing and pension markets increases substantially among those in the lower half of the income distribution. However, participation in the pension market does not increase among those at the very bottom due to the availability of the social security system, consistent with Hubbard et al. (1995). Reduced access to mortgages for the top half of the income distribution limits their trading opportunities for upgrading to larger houses, thereby reducing the intensive margin of home equity. When decomposing wealth accumulation by income into active saving through financial contracts and capital gains from asset price changes, we find that worse access to asset markets further amplifies differences in wealth accumulation. The reason is that access provides not only the opportunity to accumulate wealth through active saving, but also the opportunity to generate additional wealth growth through asset price increases. To summarize, the model with homogeneous access to asset markets overstates wealth-building opportunities for low-income households and understates them for high-income agents relative to the data.

Our model provides a new explanation for why poor households save little. It is not primarily because they do not want to save more; rather, it is because they do not have the opportunity to accumulate wealth in the asset classes most important for building wealth. This new perspective suggests an important role for recent policy proposals that aim to improve low-income households' access to asset markets. The last part of the paper evaluates one such proposal: the bipartisan *Retirement Savings for Americans Act*. According to the Act, workers without access to employer-sponsored retirement plans should become eligible for government retirement accounts with provisions to subsidize pensions for low-income workers. To capture this in the model, we assume that workers who receive job offers without access to an employer plan can choose to enroll in a government plan, with government contributions that phase out with income. When simulating the policy, we find it to be successful in improving the asset market participation of lower middle-class households. We validate the macroeconomic predictions from such a policy based on historical data.

Our work lies at the intersection of three strands of the macroeconomic literature. First, we contribute to the large literature studying the sources of wealth inequality. Existing research has extended the standard incomplete markets model à la Bewley (1986), Huggett (1993), and Aiyagari (1994) through various channels to match the high degree of wealth concentration in the data.[4] Recently, return heterogeneity has been emphasized as a driver of wealth concentration at the top motivated by new empirical evidence (Fagereng et al., 2021; Bach et al., 2020)

[4]The proposed channels include preference heterogeneity (Krusell and Smith, 1998), rich income dynamics (Castañeda et al., 2003), entrepreneurial activity (Quadrini, 2000; Cagetti and De Nardi, 2006), bequests (De Nardi, 2004), medical expense risk (De Nardi et al., 2010), or non-homothetic preferences (De Nardi, 2004; Straub, 2019). We refer the reader to the excellent reviews in Cagetti and De Nardi (2008) and De Nardi and Fella (2017). Kuhn and Ríos-Rull (2025) provide an overview of the literature on inequality facts.

and theoretical findings (Benhabib et al., 2019; Hubmer et al., 2021).[5] These studies focus on accounting for wealth concentration at the very top of the wealth distribution, whereas we study the large differences in wealth holdings across all income deciles. Furthermore, we study the joint distribution of income and wealth rather than the marginal distribution of wealth (Gini coefficient, top wealth shares). There are two notable exceptions to this strand of the literature with its focus on the top of the wealth distribution. Hubbard et al. (1995) demonstrate how social welfare programs can lead to low wealth accumulation among income-poor households. Dynan et al. (2004) document large differences in saving rates along the income distribution and conjecture that "higher-income households may save more if they enjoy better access to investment opportunities". We confirm this conjecture both empirically and theoretically.[6]

Second, we contribute to the growing macroeconomic literature on saving behavior and portfolio choice that emphasizes the role of financial frictions in asset allocation (Kaplan and Violante, 2014; Kaplan et al., 2018; Luetticke, 2021; Bayer et al., 2024; Bilbiie, 2024). We follow this strand in emphasizing frictional trading in asset markets. In particular, we borrow from Luetticke (2021) and Bilbiie (2024) the modeling of trading opportunities as a stochastic arrival process. The key difference of our approach is that we introduce heterogeneity in trading opportunities by income in line with our empirical evidence. Moreover, this strand of the literature differentiates between liquid and illiquid asset aggregates, while we differentiate between the three major portfolio components associated with key financial lifetime decisions. Our more disaggregated portfolio decomposition allows us to study the role of asset market participation, which we document to be the key driver of wealth differences by income. Importantly, we share with this strand that we take financial frictions as given, without microfounding them, and study their consequences for saving behavior and wealth accumulation.

Finally, we connect to the household finance literature that emphasizes the importance of asset market participation for portfolio allocation and wealth accumulation (see the surveys by Guiso and Sodini, 2013; Gomes et al., 2021). Consistent with our empirical findings, Guiso and Sodini (2013) show that participation increases with wealth across all asset classes. Their review further reveals that most of the literature has focused on stock market participation (see, for example, Mankiw and Zeldes, 1991; Haliassos and Bertaut, 1995; Guiso et al., 2002), with proposed mechanisms centered on the demand side (i.e., the investors' willingness to enter the stock market) under the assumption that all households have equal access to financial markets. Two noteworthy exceptions focus on employer-sponsored retirement accounts. Choukhmane and de Silva (2024) documents empirically that participation in these accounts results from participation frictions rather than non-standard preferences, consistent with our approach. Furthermore, Choukhmane (2025) shows empirically that increasing participation in retirement plans significantly lowers inequality in savings, in line with the implications of our quantitative model.

The rest of the paper is structured as follows. Section 2 conducts the empirical analysis. Section 3 introduces the model. The calibration strategy is discussed in Section 4. Our main findings are presented in Section 5. Section 6 conducts the policy experiment. Section 7 concludes.

[5] Hubmer et al. (2021) show that return heterogeneity is sufficient to match wealth concentration at the top.

[6] Over the past decade, there has been renewed interest in the topic of savings rate heterogeneity, see for example Saez and Zucman (2014); Alan et al. (2015); Krueger et al. (2016); Straub (2019); Fagereng et al. (2021).

2 Empirical analysis

2.1 Our approach

We rely on the *Survey of Consumer Finances* (SCF) data as the most widely used data to study the financial situation of US households (e.g., Bhutta et al., 2020; Kuhn and Rıos-Rull, 2016). To align the data with our quantitative framework, we focus on households with household heads aged 25–65, where the head is in dependent employment or self-employed. To abstract from business-cycle variation, we pool the four survey waves from 2010, 2013, 2016, and 2019, covering data over the last decade. We describe the sample selection criteria in Appendix A.1 and demonstrate in Appendix A.2 that the main findings of this section remain robust when we combine data from all waves between 1989 and 2019.

Our empirical analysis starts from two main variables: income and wealth. Wealth is the value of all financial and real assets net of all debts, i.e., marketable net worth. Income is the sum of pre-tax wage and salary income, business income (income from a sole proprietorship or a farm), and transfer income. We exclude capital income (capital gains, dividends, and interest) from total income to avoid any mechanical relationship between income and wealth and to align the definitions in the data and the quantitative model.

We focus on home equity, retirement accounts, and business wealth as the key portfolio components on the household balance sheet and refer to their sum as *core wealth*.[7] These wealth components capture the key financial lifetime decisions: buying a home, starting a retirement plan, and becoming an entrepreneur. On average, core wealth accounts for 70% of total household wealth in our sample.

The goal of our paper is to study the joint distribution of income and wealth and the differences in wealth between income groups. To characterize the joint distribution of wealth and income, we group households into income deciles and compute average wealth-to-income ratios within each income decile. This allows us to control for income differences and study how wealth accumulation differs per dollar of income, akin to differences in an accumulated savings rate.

By averaging core wealth within each income decile, our analysis abstracts from variation in wealth within income groups. To assess the importance of this simplification, we apply the decomposition method of Bhattacharya and Mahalanobis (1967), which breaks the Gini coefficient for total wealth into between-group, within-group, and residual components, and find that between-group inequality accounts for 83% of the inequality in core wealth, which strongly supports our focus on between-group differences.[8]

Figure 1(a) plots wealth-to-income ratios for total wealth (solid-circle line) and core wealth (solid-triangle lines) by income deciles. Wealth-to-income ratios for total wealth rise linearly up to the 8^{th} income decile and exhibit a steeper increase towards the top of the income distribu-

[7]Home equity is the net position of housing assets net of mortgage or related debt. Retirement wealth mostly comprises employer-sponsored retirement plans (64%) and individual retirement accounts and Keogh plans (30%), which are accounts workers typically opened with their previous employers.

[8]Heikkuri and Schief (2025) develop a novel decomposition into within- and between-group inequality only. Using their method, between-group inequality accounts for 51% of wealth inequality, thus still accounting for the larger part of overall inequality.

tion. In the first income decile, wealth corresponds to approximately annual income, whereas households in the top income decile hold wealth that is five times their annual income, so a top-income decile household accumulates four more dollars of wealth per dollar of income than a household in the bottom-income decile. This observation reinstates the well-established fact that wealth is significantly more concentrated than income.

Figure 1: Total wealth, core wealth, and portfolio components.

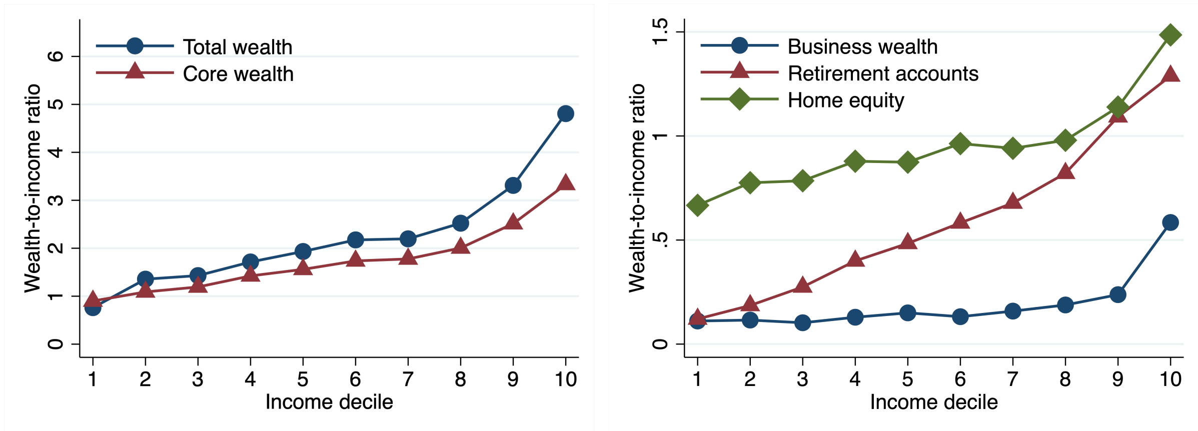

(a) Total wealth and core wealth

(b) Portfolio components

Notes: Left panel shows total wealth and core wealth relative to household income across income deciles. Right panel shows the decomposition of core wealth from the left panel into its three components. Wealth-to-income ratios are computed at the household level and then averaged within income deciles. Core wealth is the sum of business wealth, home equity, and retirement accounts. Total wealth comprises core wealth plus stocks, bonds, mutual funds, liquid assets, and remaining assets. The value of each asset component is computed net of debts. See Appendix A.1 for the sample selection criteria.

Looking at wealth-to-income ratios for core wealth in Figure 1(a) reveals two main observations. First, wealth-to-income ratios for core wealth lie below those for total wealth, consistent with core wealth accounting on average for 70% of total wealth. This share, however, varies with income: for the bottom 90% of households, it exceeds 75%, while it falls slightly below the average to 68% for the top 10%. This underscores the central role of core wealth for the vast majority of households.[9] Second, core and total wealth increase in parallel over most of the income distribution, indicating that both wealth measures display a similar pattern in their joint distribution with income. Hence, understanding the unequal distribution of core wealth by income will be highly informative about the drivers of unequal distribution of total wealth and income. In the remainder of the paper, we will, therefore, focus on core wealth and its components.

2.2 Asset market participation and wealth of participants

Figure 1(b) plots the wealth-to-income ratios separately for each component of core wealth. We observe a consistent picture across all asset classes with the same qualitative pattern as for total wealth: An increase that is close to linear up to the 8^{th} decile and a steeper increase in the two top deciles. Hence, each asset class contributes to more wealth accumulation by income and a stronger concentration of wealth compared to income in the population. The most notable difference is the very steep increase in business wealth in the top income decile, which makes

[9] Appendix A.3 offers detailed results on the portfolio composition by income.

business wealth a key driver of wealth concentration at the very top of the income distribution, as previously documented by Quadrini (2000) and Cagetti and De Nardi (2006).

The observed differences in wealth-to-income ratios can stem from differences in asset market participation or invested amounts among participating households. In the next step, we decompose the total variation in wealth-to-income ratios into the participation rate for each asset (extensive margin) and the asset value held by participants (intensive margin).[10] Denoting by ω_a^d the wealth-to-income ratio in decile d for asset class a, we have:

$$\omega_a^d = \mu_d^a \times \tilde{\omega}_a^d, \tag{1}$$

where μ_d^a is the population share of households holding positive amounts of asset a in income decile d and $\tilde{\omega}_a^d$ is the wealth-to-income ratio, conditional on holding the asset.

Figure 2: Extensive and intensive margin variation in core wealth components in the SCF.

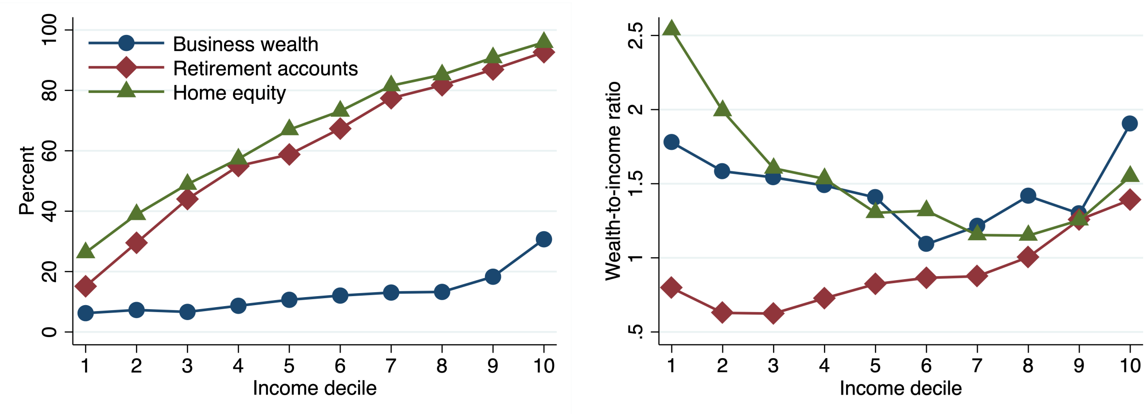

(a) Extensive margin (b) Intensive margin

Notes: Left panel: the participation rate for each asset class by income deciles. Right panel: wealth-to-income ratios for each asset class by income deciles, conditional on participation in the respective asset market.

Figure 2 presents the decomposition result. Consider first the extensive margin variation μ_d^a in Panel (a). The figure shows that participation rates in each asset component rise strongly with income, meaning that the share of households investing in home equity, retirement funds, or business wealth increases with income. The figure reveals substantial variation in participation rates by income, with a 6-fold increase in participation for retirement accounts and business wealth and a 4-fold increase for home equity. Notably, these rising patterns at the extensive margin resemble the increasing pattern in the wealth-to-income ratio in Figure 1(b).

These results are in stark contrast to the patterns of the intensive margin variation $\tilde{w}_a^d$ in Figure 2(b). The intensive margin of business equity is approximately constant along the income distribution, and it increases only slightly for retirement accounts. Strikingly, the ratio of home equity to income even declines with income.

Our findings underscore the crucial role of asset market participation in explaining differences in wealth accumulation across the income distribution. They extend the results in Guiso and

[10] We assume that a household owns an asset if the value of the asset held is strictly positive, but our results are robust to using positive cutoff values close to zero.

Sodini (2013), who emphasized the importance of differences in asset market participation by wealth for the observed wealth differences. Our results show that once a household has invested in these assets, asset accumulation relative to income varies little with income. Conditional on participation, the observed intensive margin variation implies that wealth inequality would be similar to, or even less than, income inequality.

To quantify the contribution of the extensive and intensive margins to wealth accumulation, we construct wealth-to-income ratios under two counterfactual scenarios. In the *fixed extensive margin* scenario, we set the extensive margins of the three core wealth components in each income decile at their sample average and use the observed intensive margins from Figure 2(b). We then apply eq. (1) to calculate the counterfactual wealth-to-income ratios. Reversely, the *fixed intensive margin* scenario sets the intensive margins for each asset class constant at their sample averages and uses the observed extensive margins from Figure 2(a).

Figure 3 presents the results of the two counterfactual scenarios together with the original wealth-to-income ratios (solid-triangle line). The results are striking. Eliminating the variation in the extensive margin (*fixed extensive margin*), we obtain wealth-to-income ratios that decrease up to the 7^{th} income decile, in contrast to the increasing ratios in the data. In addition, wealth-to-income ratios in the top decile are below wealth-to-income ratios in the bottom decile. By contrast, the effect of removing the intensive margin variation (*fixed intensive margin*) on the variation in wealth by income is modest for most of the income distribution, and we only find notable variation in the top decile of the distribution. As in the original data, wealth-to-income ratios remain strongly increasing over the entire income distribution.[11]

Figure 3: Extensive and intensive margin effects on wealth-to-income ratios.

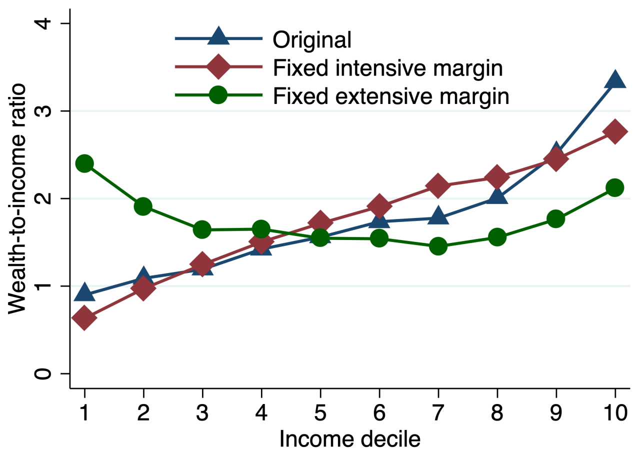

Notes: The figure shows a counterfactual decomposition of wealth-to-income differences. *Fixed extensive margin* experiment assumes the extensive margins to be constant across the income distribution and equal to the average extensive margin for a given asset class, while the intensive margins correspond to the observed values in the data. The reverse applies to the *Fixed intensive margin* experiment.

This final observation establishes our first new empirical fact on US household wealth. Variation in participation rates accounts for virtually all the differences in wealth accumulation along the income distribution. From a theory perspective, the importance of participation differences

[11]While highly instructive, this decomposition exercise is stylized because it assumes the extensive and intensive margins to remain constant by income. In the quantitative part, we will allow extensive and intensive margins to respond endogenously to changes in asset market access and allow for asset substitution effects.

scrutinizes the focus of many macroeconomic models on marginal investment decisions, in which wealth accumulation is governed by the Euler equation. While the marginal consideration captures well average wealth accumulation, it misses the substantial variation across the income distribution stemming from asset market participation decisions.

2.3 Understanding the joint distribution of wealth and income

The documented variation in asset market participation can result from differences in the demand for assets or the access to asset markets. In this subsection, we present our second new empirical finding that there are large differences in observable asset market access by income. The SCF data are instrumental for this analysis since we observe direct measures of market access for employer-sponsored retirement accounts and the housing (mortgage) market.

2.3.1 Pension market

The SCF data allow us to study market access to employer-sponsored retirement plans because households are asked a series of questions directly aimed at learning about access to this market. For this analysis, we restrict the SCF sample to households whose heads report working in dependent employment, but we keep the boundaries of the income deciles the same to maintain consistency with the preceding analysis.

Approximately 40% of household heads do not have an active employer-sponsored pension account. This is either because their employer does not offer such a plan (28%) or because they are not eligible for the plan offered by their employer (12%). The remaining 60% of workers have active employer-sponsored retirement plans. We define the *offer rate* as the fraction of employees whose employer offers a retirement plan and who are eligible to participate in the plan. Looking at workers who received an offer for an employer plan, we find that 80% of these households are enrolled in the plan. We refer to this measure as the *pick-up rate*.[12]

The solid-circle line in Figure 4(a) uncovers the large variation in offer rates along the income distribution. The offer rate increases by a factor of four from approximately 20% in the bottom income decile to more than 80% in the top decile. The solid-triangle line depicts the variation in the pick-up rates by income. We observe a higher level and a substantially lower variation in pick-up rates than in offer rates. The pick-up rate increases only by a factor of 1.7 from the bottom to the top decile.

The variation in offer rate, which is the direct measure of market access for retirement accounts, closely aligns with the variation in participation rates, previously shown in Figure 2(a). We take this finding as evidence that access heterogeneity by income is an important driver of participation rate differences by income. We will explore if this interpretation is consistent with economic theory using our model in Section 3. In the model, we will take the observed offer rates by income as given. The decision to participate and the choice of which jobs to accept will be endogenously determined, so the resulting participation rates and wealth accumulation will provide an opportunity to test our interpretation of the empirical results.

[12]This naming constitutes a slight abuse of concepts as the offer rate in the SCF data is a stock of households having received an offer rather than the share of households receiving an offer per period of time (flow rate). The same applies to pick-up rates. We will take this into account when calibrating the model to these data.

Figure 4: Access heterogeneity by income decile.

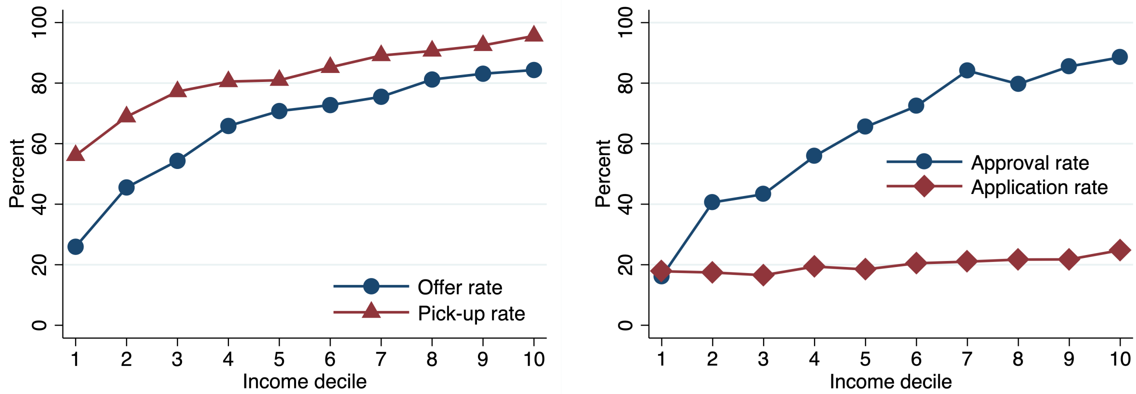

(a) Pension market (b) Housing market

Notes: Left panel: Offer rates and pick-up rates in the market for employer-sponsored retirement accounts. Offer rates are the fraction of employees whose employer offers a retirement plan and who are eligible to participate among all workers. Pick-up rates are the proportion of employees who enroll in the employer-sponsored retirement plan out of the total pool of employees who are offered a plan. Right panel: Approval rates and application rates in the housing market. An applicant is a household who submitted a loan application in the past 12 months or who did not submit an application in the past 12 months because he/she thought it would be denied. An application is treated as approved if the household receives full funding.

Following the tradition of incomplete markets literature, we will not microfound why certain trading opportunities do not exist. However, as part of the policy experiment in Section 6, we will present evidence that job-related differences are likely important drivers of workers' differential access to employer-sponsored retirement plans. One concern regarding our evidence on missing access is that it could be related to unattractive plan conditions for low-income workers. In Appendix A.5.2, we argue against this alternative explanation by documenting that there are no systematic differences in contribution and copay rates that could lead to the observed variation in participation rates.

2.3.2 Housing market

Approval of a mortgage application is the key condition for accessing the housing market. To study differences in access to the housing market, we rely again on our SCF dataset since it provides information on mortgage approval rates by income. We define a household as an applicant if it has applied for a mortgage or home-based loan or requested to refinance an existing mortgage in the past 12 months. We assign to the group of rejected applicants in addition all those households that did not apply because they anticipated being rejected as well as households who did not receive full funding.[13]

In our SCF sample, 20% of households applied for a mortgage or refinancing during the past 12 months, and, on average, two in five households got rejected or received only partial funding. Figure 4(b) unveils, however, that the average approval rate hides a lot of heterogeneity. The solid-circle line shows how approval rates vary with household income (as before, the boundaries of the income deciles are kept unchanged). The approval rate at median income is 60%,

[13] In our sample, 5% of households did not apply because they anticipated their application to be rejected.

which is very close to the average approval rate of all mortgage applications of 62%. Along the income distribution, we observe substantial heterogeneity in approval rates increasing by a factor of five from 16% in the bottom income decile to roughly 80% in the 7^{th} income decile and 90% in the top decile. The increase by income is non-linear: Approval rates rise steeply by 60pp up to the 7^{th} decile, where they almost level off and only increase by another 10pp to the top decile.[14,15] The solid-diamond line in Figure 4(b) shows the share of applicants for a mortgage. This share remains virtually constant along the income distribution at 20%.

As in the case of employer-sponsored pension plans, the described pattern in approval rates closely aligns with the increasing participation rates for home equity documented in Figure 2(a). Mortgage applications, which we take as a revealed preference measure of mortgage demand, are virtually constant along the income distribution and thus unlikely to be responsible for the strong rise in the participation rates. As in the case of pension accounts, we interpret these findings as evidence for income-dependent access driving participation rates. Again, it might be the case that unattractive mortgage conditions are an alternative driver of low participation among low-income households. In Section A.5.2, we study mortgage interest rates by income and conclude that the available evidence does not support this alternative explanation. In the model, we will take the observed approval rates by income as given, while the decision to apply for a mortgage for a given house will be endogenous.

2.4 Contractual savings flows

Core wealth comprises the main asset classes in which American households accumulate life-cycle savings. In the final step of our empirical analysis, we highlight the importance of what we term *contractual savings flows* for household wealth accumulation. Using SCF data, we document that most wealth is accumulated through such flows—specifically, mortgage principal repayments and pension plan contributions. Committing to such savings flows represents major lifetime decisions, as they involve long-term financial obligations.

To compute the household's total savings through retirement plans and mortgages, we rely on the detailed information in the SCF data on these products. For retirement savings, we use information on employer and employee contribution rates to the plans of each household member, together with information on the household members' labor incomes. Multiplying contribution rates and member-specific labor income allows us to construct the appropriate levels of saving flows. For mortgages, we consider all loans taken for home purchase, home improvements, or home repairs. We multiply the outstanding mortgage balance by the mortgage interest rate to compute interest payments and then subtract interest payments from total mortgage payments for each household in the sample.

We combine retirement and mortgage savings flows and represent them as a percentage of the total household income. We find that the household saving rate from these contractual savings flows is 3.8%. To assess the significance of these contractual savings for overall household

[14] Excluding mortgage refinancing applications reduces approval rates in the lower half of the income distribution on average by 5pp, thus strengthening our finding regarding the steep gradient in approval rates.

[15] Using the *National Survey of Mortgage Originations* for 2013–2021, we find that the vast majority of mortgage applicants (ca. 80%) apply to a single bank. Hence, the pattern of application (and approval) rates depicted in Figure 4(b) is unlikely to be contaminated by households submitting multiple applications to secure approval.

savings, we compare this constructed contractual SCF saving rate to the aggregate saving rate calculated from the National Income and Product Accounts (NIPA).

The average NIPA saving rate out of disposable income for the period from 2010 to 2019 is 6.2%. This saving rate needs to be adjusted to compare it to the constructed SCF saving rate because SCF income is total income. We adjust the NIPA saving rate using aggregate data on total and disposable household income. Specifically, we use the average ratio of disposable personal income to personal income in NIPA during this time period, which is 0.9. Using this ratio, we adjust the NIPA saving rate to obtain a comparable saving rate of 5.4%.

Comparing the SCF saving rate from contractual saving flows of 3.8% to the (total) NIPA saving rate of 5.4%, we find that the former accounts for 71% of the NIPA rate. Note that in the SCF, we only consider contractual saving flows that contribute to building home equity and retirement wealth and abstract from other savings. The lower share of 71% is therefore expected and is strikingly close to the 70% share of core wealth in total household wealth in the SCF.

The finding points to a key role of long-term financial contracts for wealth accumulation. We incorporate this mechanism in our model by assuming that pre-specified mortgage and pension contracts exist and that agents who gain access and choose to enter into these contracts are committed to executing them and build wealth through this process.

In sum, the three new empirical facts motivate our assumptions about the structure of the savings decision in the model below. Rather than focusing on savings decisions regarding the marginal dollar in always-accessible asset markets, we focus on discrete lifetime decisions regarding financial contracts in asset markets with income-dependent access. The new empirical facts provide strong support for this modeling assumption, and we will explore how a structural model with these features compares to observed wealth accumulation patterns.

2.5 Alternative explanations and sensitivity analyses

This section summarizes extensive analyses whose details we relegate to the appendix. The different analyses address potential concerns regarding the presented empirical results and provide empirical support for some of the modeling choices in Section 3. First, we provide evidence that from a life-cycle perspective, core wealth captures virtually all life-cycle accumulation dynamics. Second, we provide evidence that other dimensions of heterogeneity in the mortgage and pension market are unlikely drivers of the observed differences in asset market participation. Third, we demonstrate that the correlation with income is not just a mediator of differences in household demographics, such as education or age, but that participation, access, and income remain significantly correlated even after controlling for observable household differences. Finally, we corroborate our findings from the SCF data in alternative data sources.

Non-core wealth and life-cycle accumulation: In Appendix A.4, we document two additional facts on wealth inequality: (i) non-core wealth is mainly concentrated at the top of the income distribution and (ii) households without core wealth have, on average, also zero non-core wealth. Hence, there is no evidence that households without core wealth substitute to non-core wealth for wealth accumulation. Given the focus on life-cycle wealth accumulation and between-income-group differences, these facts support our approach of abstracting from (liquid) non-core wealth in our theory framework.

Household heterogeneity and participation: We provide evidence above that wealth accumulation, participation, and asset market access are strongly correlated with income. However, observable household characteristics such as age and education could jointly drive income, wealth, and access, so that controlling for these differences could account for the observed participation and access differences by income. In Appendix A.5.1, we provide regression results controlling for household demographics (age, race, education, marital status, sex) in addition to household income. We always find that the relationship between income and wealth, income and participation, and income and access remains economically and statistically significant. We take this finding as support for our modeling assumption that there is an important part of access heterogeneity with income. In Section 6.1, we provide additional evidence in the context of our policy experiment on access to retirement plans. We demonstrate that the job-related differences in income are even more strongly correlated with access compared to the raw correlation of income and access. This evidence suggests that it is the "labor market situation" of households that plays an important role as a driver of access heterogeneity.

Alternative explanations: We have provided evidence above that heterogeneity in asset market access is strongly correlated with differences in asset market participation. Instead of differences in access, it could be differences in household demand that drive the differences in participation rates. In Appendix A.5.2, we address this concern by studying heterogeneity in factors that could affect household demand, such as mortgage interest rates, worker contribution rates, and employer copay rates by income. If anything, we find copay rates to *decline* with income, and there is no evidence that low-income borrowers do not participate in mortgage markets because of higher interest rates. Furthermore, while median differences in mortgage rates, contribution rates, and copay rates across income groups are small, we document substantially larger heterogeneity within income groups.[16] We take this small observed variation across income groups as support for our modeling assumption of uniform financial contracts for all income groups. Accounting for the much larger within-income-group heterogeneity would allow our model to generate more within-income-group wealth inequality, thereby moving the model closer to the data in this dimension.

Validation with other data: We corroborate the results on income dependence of market access from the SCF data also in other independent, high-quality data sources. In Appendix A.5.3, we use *Home Mortgage Disclosure Act* (HMDA) data to study the relationship between mortgage approval rates and income. The HMDA data allow us to control not only for household demographics but also for detailed mortgage characteristics. We find that loan-to-value ratios have an insignificant zero effect when added to a regression for approval rates. Debt-to-income ratios have a significant effect, but the effect of income still remains statistically and economically significant. In Appendix A.5.4, we rely on publicly available data from the *National Compensation Survey* (NCS) and results from a special data request to the BLS for this survey as alternative and independent data to corroborate our results on market access by income for employer-provided pension plans. Based on these data, we corroborate that offer rates strongly increase with income.[17] As in the SCF data, pick-up rates vary less than offer rates by income.

[16] Choukhmane et al. (2024) use a novel employer–employee linked dataset and also document large residual heterogeneity in contribution rates when controlling for income.

[17] Lower variation in offer rates in the NCS data results from the fact that the SCF takes eligibility into account, while the NCS does not.

3 The model

Guided by the empirical evidence, we develop a life-cycle model of wealth accumulation centered around three financial lifetime decisions: buying a home, starting a retirement plan, and becoming an entrepreneur. The model focuses on life-cycle wealth accumulation as the most important saving motive (e.g., Gourinchas and Parker, 2002; Cagetti, 2003), and we abstract from transitory income shocks and low-return liquid assets used to accumulate precautionary savings to smooth out such transitory shocks (Kaplan and Violante, 2014; Bayer et al., 2019).

We use the new model to quantify the role of heterogeneous asset market access for differences in wealth accumulation. Consistent with the empirical findings, we focus on asset market participation (Section 2.2) and treat access to asset markets—such as mortgage approval or employer-sponsored pension plan offers—as heterogeneous and exogenous (Section 2.3). While asset market access is exogenous, the saving decision and wealth accumulation are not, as households optimally choose whether to participate in asset markets in response to the prevalence of investment opportunities. For asset market participants, wealth accumulation proceeds through contractual savings flows in fixed financial contracts (Section 2.4).

Our model also accounts for an additional consequence of asset market access: asset market participation enables households to benefit from capital gains through asset price changes. Market access, therefore, provides not only saving opportunities but also access to an income-generating technology, further amplifying differences in wealth accumulation between households with and without access.

3.1 Overview

There is a continuum of finitely-lived agents. Agents enter the model at age $j = 1$, participate in the labor market during the first J periods, and retire at age $J+1$. During retirement, the agent may die with a fixed probability $1 - \rho$ each period so that the retirement period is of stochastic length $1/(1 - \rho)$.

Each model period is split into two stages: a labor market stage and a housing investment stage. At the labor market stage, households receive job offers and business opportunities. Job offers may come with a retirement plan, so the agent needs to decide whether to accept the job offer and, conditional on accepting it, whether to enroll in a plan if offered. Business opportunities provide an alternative to dependent employment. Entrepreneurs can always set up a retirement plan for themselves when starting a new business. At the housing investment stage, agents receive housing offers and decide whether to apply for a mortgage to buy a house.

Agents accumulate wealth in retirement accounts p and home equity $qh - m$, where h denotes the size of the house, q the current house price, and m the mortgage balance. The dividend flow d of an active business is associated with business wealth b (details below). Hence, the agent's total wealth, w, is given by:

$$w = p + qh - m + b \tag{2}$$

and, therefore, matches the definition of core wealth from the empirical analysis.

3.2 Preferences

During the working stage, the agent's instantaneous utility function is additively separable in non-durable consumption c and housing services that are proportional to the house size h:

$$u(c, h) = \alpha \log(c) + (1 - \alpha) \log(\underline{h} + h), \tag{3}$$

where α is the weight on non-durable consumption and $\underline{h}$ denotes a minimum level of housing received by each agent that we introduce to abstract from a separate renting market.

3.3 Retirement

We focus on life-cycle wealth accumulation during working life. We model retirement in a reduced form, assuming that agents derive utility from their accumulated wealth and, in addition, receive a lump-sum social security benefit with present discounted value g. We include social security benefits in our analysis since previous studies have demonstrated that they create negative incentives for wealth accumulation among low-income households (Hubbard et al., 1995). To derive the utility from wealth in retirement, we solve a deterministic consumption-saving problem where consumption choices depend on total wealth only, agents have per-period utility $u(c) = \log(c)$ and face a fixed probability ρ to survive to the next period. The solution to this problem is to consume a constant fraction $1 - \beta\rho$ of total wealth $w + g$ each period, where $w = p + qh - m + b$ is accumulated lifetime wealth. This results in the value function of wealth at the entry of retirement

$$v(w) = \frac{\log(w + g)}{1 - \beta\rho} - \frac{\log(1 - \beta\rho)}{1 - \beta\rho}. \tag{4}$$

We will use this value function during retirement in the recursive representation of the model.

3.4 Labor market stage

At the labor market stage, the agent enters as a worker or entrepreneur. The worker's labor market state is denoted by l. A worker in a labor market state l supplies labor inelastically and earns pre-tax labor income $l > 0$. If an agent enters as an entrepreneur, she owns a business whose market value is denoted by $b > 0$. We model business ideas as seeds to Lucas trees that, if realized, provide per-period dividends d. Motivated by the empirical evidence of roughly constant business wealth-to-income ratio at the intensive margin (Figure 2(b)), we assume that business wealth b corresponds to a constant multiplier Ψ of business income d:[18]

$$b = \Psi d. \tag{5}$$

At the beginning of the labor market stage, agents may receive new job offers for the next period. We denote the agent's current labor market state by (l, d), the job offers and business opportunities they receive by $(\tilde{l}, \tilde{d})$, and the chosen employment state for the next period by

[18] Business wealth corresponds to the firm's equity, and we implicitly assume that all business equity is intangible capital. The simplifying assumption that agents do not invest physical capital in starting a firm can be viewed as an outcome of a well-functioning venture capital market where entrepreneurs with an idea always receives financing.

(l', d'), and we will follow the same convention for other variables. To keep the model tractable, we assume that a worker never holds business wealth and a business owner does not receive labor income. Therefore, our convention is to set $b = d = 0$ if the agent is a worker $(l > 0)$ and to set $l = 0$ if the agent is a business owner $(b = \Psi d > 0)$.

In each period, a worker receives a job offer with probability P_W and a business opportunity with probability P_E. The probability of receiving neither a job offer nor a business opportunity is $1 - P_W - P_E$, in which case the agent remains in her current labor market state. The worker's idiosyncratic job offer $\tilde{l} > 0$ follows an AR(1) process (in logs):

$$\log(\tilde{l}) \quad = \quad \rho_W \log(l) + \varepsilon \quad \text{with } \varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2). \tag{6}$$

Conditional on receiving an offer $\tilde{l}$, the worker faces an exogenous probability $\pi_P(\tilde{l})$ of being offered an employer-sponsored retirement plan, and $\tilde{\psi}$ is an indicator function equal to one if the job offer comes with an option to open a pension account and zero otherwise.

Workers decide whether to accept the job offer, internalizing that the chosen labor market state will affect both the probability of being granted a mortgage during the housing investment stage later on and the labor market dynamics in the subsequent periods according to eq. (6). Conditional on accepting the job offer with an option to set up a pension account, a worker must choose whether to enroll in the plan. If she rejects the job offer, she remains in her current labor market state with $l' = l$ and $\psi' = \psi$, where ψ indicates whether she is enrolled in a plan with her current employer. We formulate the decision problem in Section 3.8 recursively.

With probability P_E, a worker receives a business opportunity with a dividend flow $\tilde{d}$. For simplicity, we assume that the business opportunity arriving to workers is an i.i.d. random variable drawn from a uniform distribution and is, therefore, uncorrelated with the worker's current job. Workers who receive a business opportunity must decide whether to realize it.

Entrepreneurs derive a dividend income d each period and consume the value of their business b during retirement according to eq. (4). Once agents decide to become entrepreneurs, they cannot switch back to being workers for the remainder of their working lives. Entrepreneurs can grow their business by receiving a new business opportunity $\tilde{d}$ each period, which follows a stochastic process equivalent to eq. (6). Upon receiving a new business opportunity, entrepreneurs must decide whether to realize or forgo it. The private pension market for entrepreneurs is assumed to be frictionless, i.e., all business ideas arrive with an opportunity to open a retirement account, and we allow business owners to choose whether they want to set up a private pension account whenever they accept a new business opportunity.

3.5 Retirement accounts

Workers and entrepreneurs receive opportunities to invest in pension accounts. While business owners can always access the pension market when receiving a new business opportunity, access to employer-provided retirement accounts is frictional. Consistent with our empirical analysis, we assume that the probability of being offered an employer-sponsored retirement account, $\pi_P(\tilde{l})$, depends exogenously on the worker's labor market state $\tilde{l}$. We use the data from Figure 4(a) to parameterize this probability schedule (see Section 4 for details).

While the supply of employer-sponsored plans is exogenous, our model generates an endogenous demand for these assets. Retirement plans are attractive for life-cycle wealth accumulation for three reasons. First, they are subject to preferential tax treatment as contributions are made from pre-tax labor income. Second, the employer matches a portion of the employee's contributions. Third, pension wealth is typically invested in financial markets in a combination of bonds and stocks that yield interest income, dividends, and capital gains. Our framework incorporates all three features, making retirement plans an important means of wealth accumulation.

We model the retirement plan as a pre-specified financial product where workers commit to pay a fixed share χ_W of their labor income l into the account each period as long as they work in the same job. The employer matches the employee's contribution at a fixed rate σ_P, and workers can deduct the amount $\chi_W l$ from their taxable income. Each period, the account holder receives a rate of return r_P on her pension account balance p that comprises interest income, dividends, and capital gains that we assume are directly accumulated in the account. Therefore, the law of motion for pension wealth reads:

$$p' = (1 + r_P)p + \psi'(1 + \sigma_P)\chi_W l'. \tag{7}$$

If a worker chooses to stay in her current job with an active pension plan ($\psi' = \psi = 1$ and $l' = l$), she continues to pay into the retirement account. Employer-provided pension plans are fully portable across jobs, so when workers move from one job to another, they carry their accumulated retirement savings with them and continue paying into the account as long as they again enroll in a pension plan in a new job. If the new job does not offer a plan or the worker does not enroll ($\psi' = 0$), the worker keeps the existing account so that the next period's balance equals $p' = (1 + r_P)p$. Similarly, if the worker becomes an entrepreneur, she keeps her account and may choose to continue paying into it as an entrepreneur.

Entrepreneurs can also accumulate savings in a pension account, and we allow them to open an account whenever they realize a new business opportunity. Like employer-provided plans, pension accounts are pre-specified financial products that require paying each period a fixed share χ_E of the dividend payments d. Similar to eq. (7), the law of motion for pension wealth for business owners reads:

$$p' = (1 + r_P)p + \psi'\chi_E d'. \tag{8}$$

3.6 Housing investment stage

At the housing investment stage, agents receive investment opportunities for housing. We model the housing market as a search market where agents receive each period a random opportunity to buy a house of size $\tilde{h} \in \mathcal{H}$ at current price q. Housing offers follow a log-normal distribution with mean μ_h and variance σ_h^2. We introduce an additional state $\tilde{h} = 0$ in the support $\mathcal{H}$, which we interpret as an opportunity to move to rental housing and consume housing service $\underline{h}$. We denote the distribution of housing offers over the discretized support $\mathcal{H}$ by $\lambda(\tilde{h})$. A housing offer is an i.i.d. random variable across agents and time. We allow the house price q to change over time to capture capital gains in the housing market that have been a characteristic feature of housing markets over recent decades. For tractability, we assume that households have perfect foresight for house prices.[19]

Agents can always reject a housing offer and remain in their current house. Agents who prefer to downsize and draw an offer $\tilde{h} < h$ always get their mortgage application approved. If the agent is a first-time home buyer ($h = 0$) or decides to increase the value of existing housing and draws a corresponding offer ($\tilde{h} > h$), she must take out a mortgage $m \geq 0$ to finance additional housing. Consistent with our empirical analysis, we assume that the probability of being approved for a mortgage, denoted by $\pi_H(y')$, depends on the household's income $y' \in \{l', d'\}$. Workers and entrepreneurs face the same frictions in accessing the housing market, conditional on their income. To parameterize the probability schedule $\pi_H(y')$, we will directly use the approval rates for mortgage applications from the empirical analysis (Figure 4(b)). The probability of being granted a mortgage is independent of the housing value $q\tilde{h}$.[20]

While access to the mortgage market is frictional and exogenous, two forces generate endogenous demand for housing in our model. First, housing provides agents with a utility flow during the working stage according to eq. (3). Second, the accumulation of home equity $qh - m$ contributes to life-cycle savings as agents consume their accumulated wealth during retirement (eq. 4). Hence, housing serves a dual role as a consumption and investment good.

Like retirement plans, mortgages arrive as pre-specified financial contracts. A contract has a fixed interest r_M and requires each period a fixed repayment share κ of the initial mortgage amount $\underline{m}$ at the time of purchase. The law of motion for mortgages reads:

$$m' = (1 + r_M)m + q(h' - h) - \mathbb{1}_{m>0}\kappa\underline{m}, \tag{9}$$

where $\mathbb{1}_{m>0}$ is an indicator function that is one if the current mortgage balance m is positive. A mortgage balance of zero ($m = 0$) applies to non-participants in the housing market, first-time home buyers, for whom it also holds that $h = 0$, as well as to outright homeowners with $h > 0$ who have already repaid their mortgage. Although there is no downpayment for first-time home buyers, most home buyers in our model will upgrade their homes ($h' > h$) and make downpayments by using their accumulated home equity ($qh - m$). For a new house, only additional housing investments, $q(h' - h)$, will be financed through a mortgage.

All agents with outstanding mortgage debt ($m > 0$) have to pay interest r_M on their debt balance m and make an amortization payment $\kappa\underline{m}$, where $\underline{m}$ is the initial mortgage balance when purchasing the house, which becomes an additional state variable in the household problem. As there are no unexpected income declines in the model, agents have no incentive to default on mortgage contracts, so we abstract from a default decision. Finally, mortgages constitute long-term financial commitments, a fact that limits the demand for housing.

3.7 Budget constraints

An agent's taxable income is her market income $y \in \{l, d\}$ net of mortgage interest payments and contributions to retirement plans that can be made out of pre-tax income and reduce taxable

[19] There is a one-to-one mapping from household age to the house price at that age. Hence, age is a sufficient state variable for the household to know the current house price q.

[20] Appendix Table 11 shows that there is no significant effect of loan-to-value ratios on approval rates in the HMDA data once we control for income and household characteristics, which is why we abstract from this constraint.

income. Taking into account our convention that $d = 0$ if $l > 0$ and vice versa, the taxable income $\bar{y}$ becomes

$$\bar{y} = \underbrace{(l+d)}_{\text{market income}} - \underbrace{\psi(\chi_W l + \chi_E d)}_{\text{retirement contributions}} - \underbrace{\mathbb{1}_{m>0}(r_M m)}_{\text{mortgage interest payments}} . \tag{10}$$

Only agents who have an active pension plan ($\psi = 1$) contribute to their retirement account, and only agents who have a positive mortgage balance ($\mathbb{1}_{m>0} = 1$) make mortgage interest payments. Conditional on the current job and housing, the agent's non-durable consumption in the current period equals the after-tax income net of pension contributions and amortization payments in case of a positive mortgage balance:

$$c = y - T(\bar{y}) - \psi(\chi_W l + \chi_E d) - \mathbb{1}_{m>0}(\kappa \underline{m}), \tag{11}$$

where $T(\bar{y})$ denotes the income tax function specified in the calibration section.

Hence, consumption corresponds to the income net of the financial contracts and taxes and follows hand-to-mouth dynamics, with few poor hand-to-mouth households and a majority of wealthy hand-to-mouth households, in line with the empirical evidence in Kaplan et al. (2014). Note that the wealth dynamics of wealthy hand-to-mouth agents in our model differ from Kaplan and Violante (2014). Whereas agents in their model build up liquid wealth to deposit it infrequently into their illiquid account and follow hand-to-mouth behavior if their wealth is tied up in illiquid assets, wealth accumulation in our model happens constantly through long-term financial contracts (Section 2.4).

3.8 Dynamic programming problem

We summarize the agent's joint labor market and financial state at the beginning of a period by her age j and a state vector $\boldsymbol{x} \equiv (l, d, \psi, p, h, m, \underline{m})$, where (l, d) denotes the labor market state, ψ the indicator function whether the agent has an active pension plan, p the pension account balance, h the housing stock, m the mortgage balance, and $\underline{m}$ the initial mortgage balance when purchasing the house.

We describe the dynamic programming problem proceeding backward. Agents enter retirement at age $J + 1$. Their value function in retirement is $v(w)$ derived in Section 3.3. During working life, each period is divided into two consecutive stages: the labor market stage and the housing investment stage. Proceeding backward, we start with the housing investment stage, where the agent receives housing offers $\tilde{h}$ and decides whether to apply for a mortgage or not.

To save on notation, let the state vector $\boldsymbol{z}' = (l', d', \psi', p')$ summarize the outcome of the labor market stage of the current period. The decision about moving to a smaller house $\tilde{h} \leq h$ does not involve risk about mortgage approval, so if the agent takes the offer for a smaller house, her value function is given by $V_{j+1}(\boldsymbol{z}', \tilde{h}, m', m')$, and if she declines it, the value is given by $V_{j+1}(\boldsymbol{z}', h, m', \underline{m})$, where the last state always denotes the initial mortgage balance $\underline{m}$. If the agent applies for a mortgage for the offer $\tilde{h} > h$ and receives approval for financing, which happens with probability $\pi_H(y')$, her value function is given by $V_{j+1}(\boldsymbol{z}', \tilde{h}, m', m')$. If she does not apply for a mortgage or the offer is not approved, she remains in the current house, and her continuation value is $V_{j+1}(\boldsymbol{z}', h, m', \underline{m})$. In both cases, the mortgage balance m' follows the law of motion in eq. (9).

At the beginning of the housing investment stage, agents form expectations over housing offers, and their expected value can be summarized as follows:

$$\begin{aligned}\bar{V}_j(\boldsymbol{z}', h, m, \underline{m}) \quad = \quad & \beta \sum_{\tilde{h}>h} \lambda(\tilde{h}) \Big(\pi_H(y') \max \Big\{ V_{j+1}(\boldsymbol{z}', \tilde{h}, m', m'), V_{j+1}(\boldsymbol{z}', h, m', \underline{m}) \Big\} \\ & + (1-\pi_H(y')) V_{j+1}(\boldsymbol{z}', h, m', \underline{m}) \Big) + \beta \sum_{\tilde{h}\leq h} \lambda(\tilde{h}) \max \Big\{ V_{j+1}(\boldsymbol{z}', \tilde{h}, m', m'), V_{j+1}(\boldsymbol{z}', h, m', \underline{m}) \Big\}.\end{aligned} \tag{12}$$

Observe that if the agent does not apply for a mortgage for a housing offer $\tilde{h} > h$, the expression in round brackets boils down to $V_{j+1}(\boldsymbol{z}', h, m', \underline{m})$, which is the value of staying in house h.

At the labor market stage, the agent may receive a job offer or a business opportunity $(\tilde{l}, \tilde{d})$. If an agent receives a job offer, she also learns if it comes with a retirement account, which is denoted by $\tilde{\psi} \in \{0, 1\}$. For business owners, access to the pension market is frictionless, always leading to $\tilde{\psi} = 1$, i.e., an opportunity to set up a pension plan (so we will omit it to simplify notation). Note that agents can choose to enroll in a pension plan (if offered one) only when they switch jobs or start a new business.

The agent's decision problem in the labor market stage for different combinations of offers is

$$V_j^{WP}(\boldsymbol{x}; \tilde{l}, 1) \quad = \quad \max\{\bar{V}_j(\boldsymbol{x}), \max\{\bar{V}_j(\tilde{l}, \tilde{d}, 0, p, h, m, \underline{m}), \bar{V}_j(\tilde{l}, \tilde{d}, 1, p, h, m, \underline{m})\}\}, \tag{13}$$

$$V_j^{WN}(\boldsymbol{x}; \tilde{l}, 0) \quad = \quad \max\{\bar{V}_j(\boldsymbol{x}), \bar{V}_j(\tilde{l}, \tilde{d}, 0, p, h, m, \underline{m})\}, \tag{14}$$

$$V_j^{E}(\boldsymbol{x}; \tilde{d}) \quad = \quad \max\{\bar{V}_j(\boldsymbol{x}), \max\{\bar{V}_j(\tilde{l}, \tilde{d}, 0, p, h, m, \underline{m}), \bar{V}_j(\tilde{l}, \tilde{d}, 1, p, h, m, \underline{m})\}\}, \tag{15}$$

where superscripts indicate the occupation (W for worker and E for entrepreneur) and the opportunity to open a pension account (P when a retirement plan is available and N if not). Specifically, the first line is the worker's decision problem with a job offer $\tilde{l}$ that includes a retirement account ($\tilde{\psi} = 1$). $\bar{V}_j(\boldsymbol{x})$ stands for the agent's utility at the beginning of the housing stage, introduced above, if the agent does not change the labor market state, leading to $l = l'$, $d = d'$, $\psi = \psi'$, and p' follows from eq. (7)–(8). $\bar{V}_j(\tilde{l}, \tilde{d}, 0, p, h, m, \underline{m})$ denotes the continuation value from accepting the job offer but rejecting the pension plan offer, while $\bar{V}_j(\tilde{l}, \tilde{d}, 1, p, h, m, \underline{m})$ denotes the value from accepting the job offer and enrolling in the plan. The second line is the worker's decision problem without the opportunity to invest in a retirement plan ($\tilde{\psi} = 0$). The third line corresponds to the case when the agent receives a business opportunity $\tilde{d}$ and can set up a retirement plan if preferred.

Combining eq. (13)– (14), we further define

$$V_j^W(\boldsymbol{x}; \tilde{l}) = \pi_P(\tilde{l}) V_j^{WP}(\boldsymbol{x}; \tilde{l}, 1) + (1 - \pi_P(\tilde{l})) V_j^{WN}(\boldsymbol{x}; \tilde{l}, 0)$$

as the expected value of receiving a job offer $\tilde{l}$, where $\pi_P(\tilde{l})$ is the probability of being offered an employer-sponsored retirement account. Combining the case of receiving a job offer, a business opportunity, and the event of receiving no offer, occurring with probabilities P_W, P_E, and $1 -$

$P_W - P_E$, respectively, the value function at the beginning of the labor market stage is

$$V_j(\boldsymbol{x}) = u(c,h) + P_W \mathbb{E}[V_j^W(\boldsymbol{x};\tilde{l}) \mid l] + P_E \mathbb{E}[V_j^E(\boldsymbol{x};\tilde{d})] + (1 - P_W - P_E)\bar{V}_j(\boldsymbol{x}).$$

Since the business opportunity for workers is an i.i.d. random variable, the second expectation operator is unconditional. Consumption follows from the budget constraint in eq. (11) as the residual of the decision from the previous period and the current labor market status. Finally, since business owners do not return to dependent employment by assumption, we can use $V_j^E(\boldsymbol{x};\tilde{d})$ defined in eq. (15) to express their value function at the beginning of the period as

$$V_j(\boldsymbol{x}) = u(c,h) + \mathbb{E}\left[V_j^E(\boldsymbol{x};\tilde{d}) \mid d\right]. \tag{16}$$

4 Calibration

We calibrate the model to the SCF data from the empirical analysis in Section 2. For asset returns, we rely on data from Jordà et al. (2019). We set some parameters outside the model and calibrate the remaining parameters inside the model. Although all of the latter parameters must be calibrated jointly, we provide an intuitive discussion about the relationship of each parameter to a corresponding data moment. Table 1 lists all model parameters and reports the corresponding source or the empirical target.

To initialize the distribution of agents, we proceed as follows. All agents are assumed to enter the model at age 20. The initial conditions for these agents are parameterized based on the subsample of households aged 20–24 from the SCF data. When comparing the model to the data, we discard this age group to avoid the influence of the initial distribution, focusing instead on the age group 25–65, consistent with the sample selection in the empirical section.

At age 20, agents enter the model in the lowest labor market state and without any pension wealth, but there is a share of these agents who work in jobs offering a retirement account that we set to 30%. The share of entrants with positive home equity is 17%. We set their initial house value and the mortgage balance such that their home equity-to-income ratio, $(qh-m)/y$, is equal to 3.9 and the loan-to-value ratio, $m/(qh)$, is 53%.

A subset of the model parameters, for example, those governing the stochastic process for labor income or business ideas, is occupation-specific. To calibrate these parameters, we define a household in the SCF data as a *worker* if the household head reports working for someone else and as a *business owner* if the household head reports actively managing a business.

Demographics: The model period is one year. Agents enter at age 20 and retire at age 66. After dropping the first five years, the duration of the working life is $J = 41$ periods. The average duration of retirement is $1/(1-\rho) = 20$ years, such that the survival probability ρ is 0.95.

Labor markets: The idiosyncratic job offers $\tilde{l}$ follow an AR(1) process given by eq. (6), which we approximate using a discrete-state Markov chain. We match the variance profile of log income for workers by age from the SCF. The slope of the profile informs the dispersion of offers σ_ϵ^2, and the auto-correlation ρ_W is informed by the curvature of the profile. Importantly, the calibration of these parameters cannot be done outside the model because agents have an

Table 1: Calibrated parameters.

Parameter	Interpretation	Value	Source/Target
		Parameters set outside the model:	
Demographics:			
J	Duration of working stage	41	
ρ	Duration of retirement is $1/(1-\rho)$	0.95	
Frictional access to pension and housing markets:			
$\pi_H(y)$	Mortgage approval rates by income	Figure 4(b)	SCF
Financial contracts and returns:			
χ_W	Contribution rate for workers	0.05	SCF
r_M	Interest rate on mortgage	0.02	FRED
Ψ	Business wealth multiplier	1.5	SCF
r_P	Return on pension wealth	0.029	Jordà et al. (2019)
r_H	Return on home equity	0.005	Jordà et al. (2019)
Income taxation and Social Security:			
(τ_1, τ_2)	Income tax level and progressivity	(0.08, 0.22)	Heathcote et al. (2020)
R	Replacement rate	0.40	Brendler (2023)
		Internally calibrated parameters:	
Labor market:			
$(\rho_W, \sigma_\epsilon^2)$	AR(1) process for labor income	$(0.95, 0.06)$	Variance of labor income by age
(ρ_E, σ_η^2)	AR(1) process for bus. income	$(0.89, 0.15)$	Business wealth/labor income by age
P_W	Probability to receive job offer	0.48	Average job tenure = 7 years
P_E	Probability to receive business idea	0.02	Extensive margin for business wealth = 12.7%
Retirement accounts:			
$\pi_P(l)$	Offer rates for employer-provided plans	Schedule	See text.
β	Discount factor	0.97	Extensive margin for pension wealth = 61%
σ_P	Employer's share paid into account	0.59	Intensive margin for pension wealth = 0.9
χ_E	Contribution rate for entrepreneurs	0.08	$(1+\sigma_P)\chi_W$
Housing:			
α	Weight on non-durable consumption	0.92	Extensive margin for home equity = 67%
κ	Amortization rate	0.06	Intensive margin for home equity = 1.5
(μ_h, σ_h^2)	Log-normal distrib. of housing offers $\tilde{h}/Y$	$(1.45, 0.84)$	Mean(h/Y) = 3.0, Var(h/Y) = 11.1
$\lambda(\tilde{h}=0)$	Probability to receive rental housing	10.0%	Mortgage application rate = 20%
$\underline{h}$	Housing for renters	0.1	Lowest positive point in the housing grid $\mathcal{H}$

acceptance-rejection decision for job offers, such that the age profile of income is endogenous.[21] Appendix Figure 22(a) shows the fit to the empirical variance profile for labor income by age. Finally, the probability of receiving a job offer, P_W, is calibrated to match the average job tenure in the US labor market of 7 years (Jung and Kuhn, 2019).

The idiosyncratic process for business ideas of entrepreneurs, $\tilde{b}$, follows an AR(1) process equivalent to eq. (6) with corresponding parameters (ρ_E, σ_η^2). As for the process of labor market opportunities, we approximate it using a discrete-state Markov chain. The number of observations for business wealth in some age groups is very small in the SCF, so we target the average ratio of business wealth to income by age rather than the variance to calibrate (ρ_E, σ_η^2). This ratio is informative about the dispersion of business opportunities because it governs the decision of workers in the model whether to become entrepreneurs or stay in their current job. Appendix Figure 22(b) shows the model fit of this margin. We calibrate the arrival rate of busi-

[21] Low et al. (2010) demonstrate the importance of accounting for selection when estimating income risk due to the job change.

ness ideas P_E to match the average extensive margin for business wealth by income equal to 12.7% from the SCF sample. Note that the intensive margin of business wealth in the model is constant along the income distribution and equal to the business wealth multiplier Ψ in eq. (5). We set $\Psi = 1.5$ to match the intensive margin in the SCF data (Figure 2(b)).

Housing and pension markets: We parameterize the income-specific access rates to the mortgage market, $\pi_H(y)$, outside the model using the income-dependent approval rates for mortgage applications from the empirical analysis (Figure 4(b)). To smooth out these probabilities, we approximate it by the functional form $\alpha_0 + \alpha_1/y$. We have to proceed differently when calibrating the offer rates for employer-provided pension plans, $\pi_P(y)$. As already mentioned in the empirical analysis, the observed offer rate in the SCF data is the share of households having received an offer (stock) rather than the share of households receiving an offer per period of time (flow). For this reason, we assume that the (unobserved) flow offer rates, $\pi_P(y)$, follow the same functional form as the mortgage approval rates, and we calibrate its parameters inside the model, matching the profile of stock offer rates observed in the data (Figure 4(a)). Appendix B.2 displays the probability schedules $\pi_H(y)$ and $\pi_P(y)$ used in the model and provides an additional discussion on their specification.[22]

Housing offers $\tilde{h} \in \mathcal{H}$ follow a (discretized) log-normal distribution with mean μ_h and variance σ_h^2. We normalize housing values in the model and data by the respective average household income and calibrate the two parameters to match the mean and variance of housing (conditional on homeownership) as observed in the SCF. We add to the housing grid the lowest state equal to zero, corresponding to the case when the household receives a rental offer. We set the probability of receiving such an offer inside the model, targeting the average mortgage application rate of 20%. The subsistence level of housing, $\underline{h}$, is set to the lowest positive value in the grid for housing offers $\mathcal{H}$. We set the mortgage interest rate, r_M, at 2%.[23]

To parameterize capital gains in the housing market and returns on retirement accounts, we rely on asset price and CPI data for the United States from Jordà et al. (2019) for the time period from 1970 to 2019. In the model, we abstract from aggregate risk and follow Kaplan and Violante (2014) to adjust asset prices accordingly (see their eq. (C.5) in Online Appendix C). For house prices, we use the nominal house price growth, adjusting it for CPI inflation, and arrive at an average real annual house price growth of 0.5%. For retirement accounts, we use the 2010 to 2019 SCF data to derive the average share of retirement accounts invested in the stock market, which we find to be 46.6%. For the return on the stock market, we use dividends and stock price changes. For bonds, we use average interest rates in the bond market. We adjust these financial returns for inflation using the CPI and aggregate them using the equity share from the SCF. We find a real risk-adjusted return on retirement accounts of 2.9%.[24]

[22] We demonstrate that conditional on household income, there is little variation in offer rates for employer-sponsored pension plans across different age groups in the SCF, which is why we chose to specify offer rates in the model as a function of household income only. We provide similar evidence for mortgage approval rates.

[23] According to the Federal Reserve Economic Data (FRED), the 30-year fixed rate mortgage average in the US during 2010–2019 was 4% (see `https://fred.stlouisfed.org/series/MORTGAGE30US`). Assuming an inflation rate of 2%, we obtain the parameterized value of r_M.

[24] This value falls in the middle of the range of estimates used by Kaplan and Violante (2014), who applied a rate of 2.3% for illiquid assets and 3.5% for retirement accounts for the period 1960 to 2009.

We calibrate the remaining parameters of the housing and pension markets inside the model, targeting the average extensive and intensive margin for home equity and pension wealth (we do not target heterogeneity by income). We calibrate the relative weight on non-durable consumption, α, in eq. (3) by targeting the average extensive margin for home equity at 67%. Similarly, the amortization rate κ is calibrated to match the average intensive margin of home equity, equal to 1.5. Furthermore, the discount factor β is determined by targeting the average extensive margin for pension wealth at 61%. Finally, the average employee's share of income paid into retirement accounts is $\chi_W = 5\%$ in the SCF, and we calibrate inside the model the employer's copay rate σ_P to match the intensive margin for pension wealth equal to 0.90. Entrepreneurs are assumed to contribute the sum of the employer's and the employee's shares, i.e., $\chi_E = \chi_W(1 + \sigma_P)$, equal to ca. 8%.

Income taxation and Social Security: Following Heathcote et al. (2020), we specify the income tax function as $T(\bar{y}) = \bar{y} - \mathcal{I}(1 - \tau_1)(\bar{y}/\mathcal{I})^{1-\tau_2}$, where $\bar{y}$ is the agent's taxable income defined in eq. (10) and $\bar{y}/\mathcal{I}$ is the agent's taxable income normalized by the economy-wide average income $\mathcal{I}$. We set income tax progressivity τ_2 to 0.22 based on Heathcote et al. (2020) and the income tax level τ_1 to 0.08, which implies that the mean-income agent pays an average income tax of 8%. Retirees receive each period a constant fraction $R = 40\%$ of average income $\mathcal{I}$ as public pension benefits (Brendler, 2023). The present discounted value of all future pension benefits is then given by $g = R\mathcal{I}(1 + r_P)/(1 + r_P - \rho)$, which we use to compute the value function for retirement in eq. (4).

5 Findings

Our model focuses on lifetime financial decisions and participation in asset markets, assuming unequal access to wealth-building opportunities across income groups. The first question we address is whether this novel mechanism can replicate the empirically observed patterns of wealth accumulation. Our first step is therefore to demonstrate that the calibrated model with income-dependent asset market access closely aligns with the documented empirical facts on the joint distribution of income and wealth and asset market participation. In the second step, we conduct comparative statics to explain the model's wealth accumulation mechanism and the role of income-dependent asset market access.

5.1 Model fit

Figure 5 shows the extensive and intensive margins of home equity, pension wealth, and business equity along the income distribution in the calibrated model together with their empirical counterparts. Panels (a) and (d) show the model fit of the extensive margin and the intensive margin for home equity. The model tracks the empirical profiles of both margins closely. Importantly, it matches the strong increase at the extensive margin and the declining intensive margin in the lower half of the income distribution. Panels (b) and (e) display the model fit for the extensive and intensive margin of pension wealth. Again, the model accurately matches both margins across income groups. Panels (c) and (f) show the fit for the extensive and intensive margin for business wealth by income. Despite being a stylized model of entrepreneurial activity, the model closely matches the extensive and intensive margins of business wealth.

Figure 5: Extensive and intensive margins in the model and data.

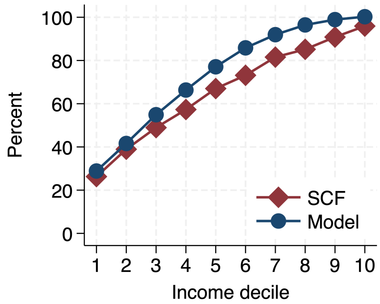

(a) Home equity: Extensive margin

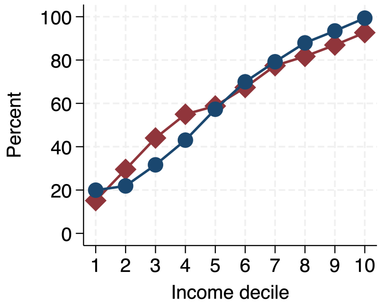

(b) Pension: Extensive margin

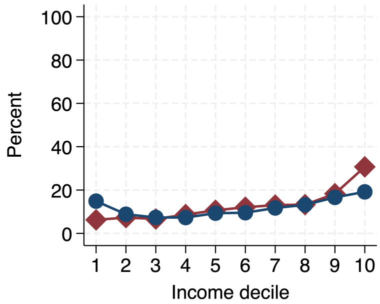

(c) Business: Extensive margin

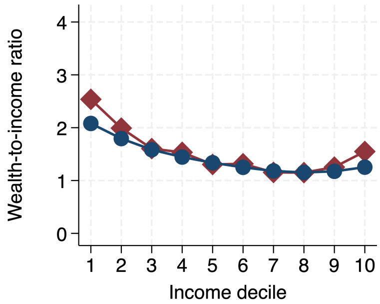

(d) Home equity: Intensive margin

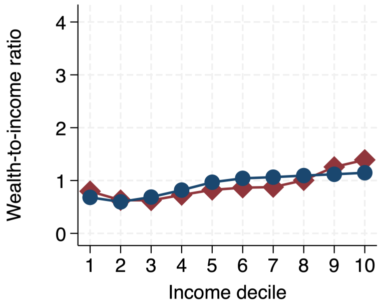

(e) Pension: Intensive margin

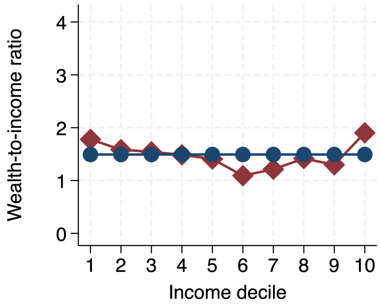

(f) Business: Intensive margin

This close fit across asset classes is not part of our calibration strategy, since we only target the average extensive and intensive margins, not their heterogeneity by income. This close fit is also not a direct consequence of exogenous access heterogeneity because participation decisions, mortgage applications, and pick-up rates for retirement plans are endogenous. Appendix C.1 further shows that the model is consistent with empirically observed constant mortgage application rates and increasing pick-up rates for retirement plans by income.

Figure 6: Wealth-to-income ratios for wealth in the model and data.

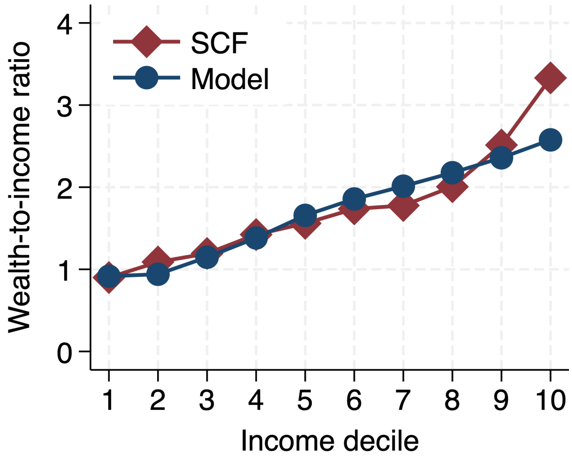

Figure 6 presents the resulting wealth-to-income ratios by income as a measure of the joint distribution of income and wealth. It shows that the model closely matches the characteristic property of increasing wealth-to-income ratios by income. The model misses only the non-linear increase of the wealth-to-income ratio towards the top decile. For a more granular decomposition, Appendix Figure 25 shows how the model fits the wealth-to-income ratios for each wealth component. Appendix Table 12 further reports the shares of each wealth component by income in the model and data. Appendix Figure 27 shows the evolution of wealth-to-income ratios by age as a non-targeted dimension of wealth accumulation, and it can be seen that the model closely aligns with the almost linear increase in wealth-to-income ratios in the data.

The overall close fit to the wealth accumulation pattern by income provides support for our wealth accumulation mechanism that combines participation decisions with contract-based wealth accumulation of participants. We take the model's ability to jointly match wealth accumulation and participation as support for this paper's view that heterogeneous asset market access is a key driver of differences in wealth accumulation along the income distribution.

Table 2: Gini index for income and wealth in the model and data.

	Model	SCF
Income	0.51	0.51
Wealth	0.69	0.77
Amplification	35%	50%

Notes: *Amplification* is a percentage change in the Gini index for net worth relative to the Gini index for income.

Table 2 compares the model fit for income and wealth inequality, as measured by the Gini index. While the income process in our model is endogenous, with workers accepting and declining job offers, we calibrate its parameters to match the increase in cross-sectional income inequality, so the model matches exactly the Gini index for income of 0.51. At the same time, wealth inequality in the model is not targeted. Looking at the Gini index for wealth in the model, we find it to be 0.69 compared to 0.77 in the data, so the model accounts for 90% of the empirical wealth inequality. Comparing the relation of income inequality to wealth inequality in the model and data, we find that wealth inequality is 35% (18pp) higher than income inequality in the model (last row of Table 2), while the amplification in the data is 50% (26pp). Thus, the model misses some of the empirical drivers of wealth inequality. We discuss this point next.

Using the decomposition approach of the wealth Gini index into between-group, within-group, and residual inequality by Bhattacharya and Mahalanobis (1967), we find a substantial contribution from between-group inequality, with Gini indices of 0.60 (model) and 0.64 (data). However, the sum of the within and residual components shows a larger discrepancy between model and data: the corresponding Gini index is 0.08 in the model versus 0.13 in the data. Applying the more recent decomposition method by Heikkuri and Schief (2025), which partitions the Gini index into only within and between components, we find that the between-group component closely matches with 0.41 (model) versus 0.39 (data), while nearly the entire gap between model and data arises from the within component: 0.28 (model) versus 0.38 (data). Therefore, the lower overall wealth inequality and the smaller amplification over income inequality, reported in Table 2, are due to lower within-income-group inequality in the model.

This conclusion of a lower within component can also be directly observed when looking at Figure 7(a) that plots the Gini coefficients for wealth within each income decile. The model repli-

Figure 7: Within-income group wealth inequality and role of capital gains.

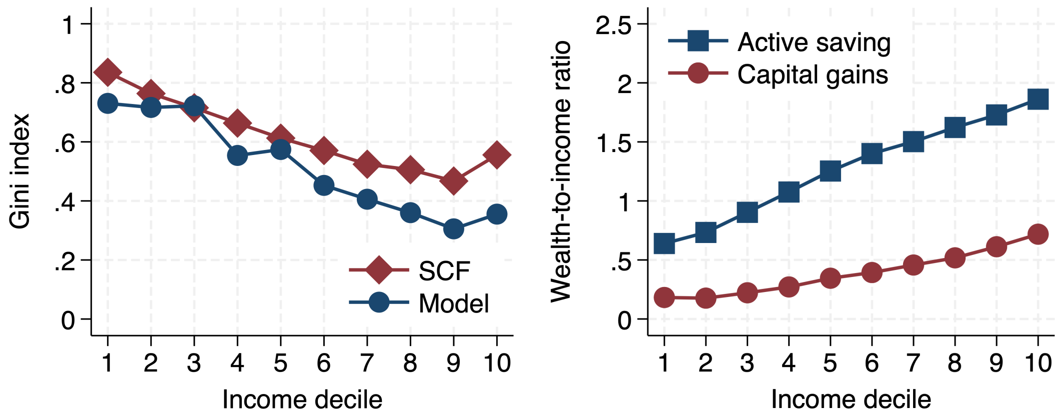

(a) Within-income inequality

(b) Capital gains and active savings

Notes: Left panel: Within-income-group wealth inequality, measured by the wealth Gini coefficient within each income decile. Right panel: Decomposition of the model-implied wealth-to-income ratio by income decile, distinguishing contributions from active savings and capital gains.

cates the empirical pattern of wealth inequality decreasing with income. Notably, the largest difference in Gini coefficients appears among the highest-income households. As discussed in the introduction, much of the existing literature focuses on explaining the high concentration of wealth at the top. The proposed mechanisms from this literature, such as heterogeneity in returns and preferences, luxurious bequest motives, and other forms of non-homothetic preferences, have in common that they introduce additional heterogeneity beyond income. In contrast, our analysis restricts household heterogeneity to income and focuses on differences in wealth accumulation across the entire distribution, making it complementary to existing work.

The presence of capital gains in the model allows us to also quantify the joint role of access to asset markets and capital gains for wealth accumulation. Figure 7(b) decomposes the wealth-to-income ratios into active saving through contributions from financial contracts and the component from asset price changes and capital gains in retirement accounts and housing following Fagereng et al. (2021). For active savings, we subtract mortgage interest payments $r_M m$ from total mortgage payments $\kappa \underline{m}$ and add employer and employee contributions to retirement accounts. For capital gains, we add, period-by-period, all returns on retirement accounts $r_P p$ and all capital gains in the housing market $(q' - q)h$ up to the household's current age. The sum of active savings and capital gains corresponds to the total household wealth, previously shown in Figure 6.

We find that capital gains account for a sizable fraction of accumulated wealth. For each dollar of accumulated wealth through active savings, households in the top decile have 40 cents of accumulated wealth through capital gains. The fact that capital gains are an important driver of wealth accumulation in our model is consistent with the empirical evidence for the United States.[25] Comparing the capital gains of the bottom decile with those of the top decile, we

[25] Kuhn et al. (2020) find that for the wealthiest 10% of US households, 40% of wealth growth between 1971 and 2007 was due to changes in asset prices. Fagereng et al. (2021) report even larger effects of asset price increases on wealth accumulation for Norway.

Table 3: Saving rates by current income in the model and data.

	Q1	Q2	Q3	Q4	Q5
Data	0.014	0.090	0.111	0.173	0.236
	(0.019)	(0.029)	(0.032)	(0.027)	(0.040)
Model	0.008	0.085	0.141	0.159	0.173

Notes: Row *Data* reprints the estimation results of median regressions of saving rate on current income from Dynan et al. (2004) (their Table 3, second column). The authors define the saving rate variable as the change in real net worth in the SCF. The sample is restricted to households aged 30–59, so we redefine income quintiles in the model based on the trimmed sample. We reproduce standard errors from their regression in brackets below each coefficient estimate. We display only the estimates for the income quintile dummies, omitting the results for other regressors. Row *Model* shows the estimation results from the same regression run inside our model.

find that households in the bottom decile have accumulated capital gains equal to 15% of their annual income, while those in the top decile have accumulated capital gains equal to 75% of their annual income, five times as much.

Through the lens of our model, the reason for larger capital gains at the top is twofold. First, these households invest higher amounts relative to their incomes. Second, they have been participants in the asset markets for longer. Thus, market access not only enables wealth accumulation but also generates additional wealth through capital gains. In this way, capital gains for market participants amplify the already existing unequal opportunities for wealth building.

Finally, Table 3 compares saving rates by income in the model with the empirical evidence from Dynan et al. (2004), where we generate model-simulated data to conduct the counterpart to their empirical estimation. We see that the model is consistent with the empirical fact of very low saving rates at the bottom of the income distribution and strongly increasing saving rates with income. Although the model saving rates tend to be lower than the point estimates in each income quintile, they are not statistically significantly different from their empirical counterparts. The largest discrepancy is the increase in saving rates between the 4^{th} and 5^{th} quintiles. The model underestimates the 6.3pp difference between these two quintiles, only generating a 1.4pp difference. It should be noted, however, that Dynan et al. (2004) include non-core wealth in their estimation. As we have documented in the empirical section (Figure 1), there is a widening gap between total and core wealth at the top of the income distribution where most non-core wealth is held, such that the increase in saving rates between income groups is naturally higher at the top income quintiles.

Taking this into account, the model provides a close fit of the empirical differences in saving rates by income and, therefore, offers a new interpretation of these differences. Through the lens of the model, the variations in saving rates are not primarily due to differences in the proportion of income saved but rather due to differences in the share of households that save. These differences, in turn, arise from variations in market access rather than a differing demand for wealth accumulation. The idea that disparities in asset market access can significantly influence observed differences in saving rates has been proposed in the discussion by Dynan et al. (2004, p. 436). Our paper provides a conceptual framework that supports this conjecture.

5.2 Income-dependent market access and wealth accumulation

In the next step, we apply the calibrated framework to explain the model's wealth accumulation mechanism and to quantify the impact of income-dependent market access. In our experiment, we change market access so that agents across the entire income distribution face the same market access in the mortgage and pension markets. Such homogeneous access is the prevailing assumption in existing models of consumption-saving behavior. Hence, our experiment quantifies the differences in wealth accumulation due to heterogeneity in market access.[26]

Figure 8: Financial frictions in the baseline model and homogeneous access experiment.

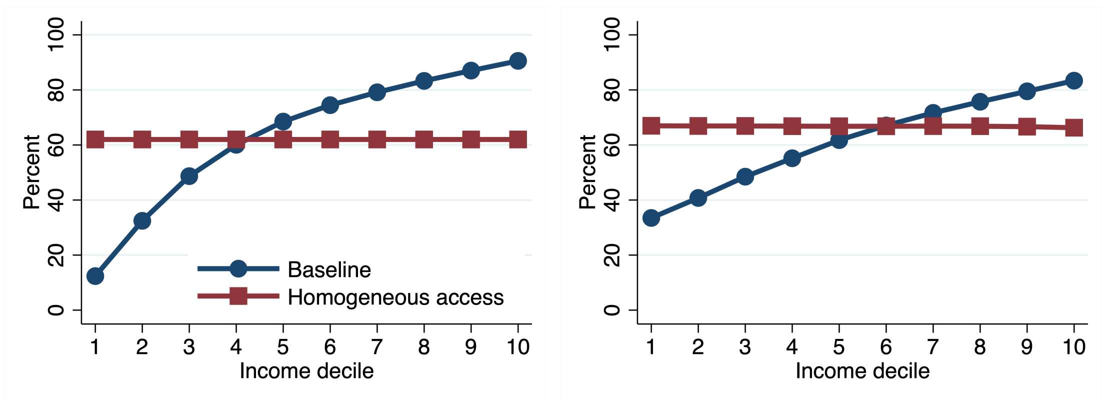

(a) Mortgage approval rates $\pi_H(y)$

(b) Offer rates $\pi_P(y)$

Notes: The figure shows the exogenous change in financial frictions in the experiment compared to the baseline model. The homogeneous access experiment sets both the mortgage approval rates, $\pi_H(y)$, and the offer rates for employer-sponsored retirement plans, $\pi_P(y)$, to their respective economy-wide averages (62% and 64%) for all income groups.

In the experiment, we set the mortgage approval rate, $\pi_H(y)$, and the offer rate for employer-sponsored accounts, $\pi_P(y)$, to their respective economy-wide averages, 62% and 64% (Figure 8). This change implies that market access to the mortgage market improves for agents with below-median incomes and worsens for those with above-median incomes. The change in access rates is the strongest for the lowest-income households, for whom mortgage approval rates jump from 15% to 62%. The change at the top is smaller, with approval rates decreasing by one-third from 90% to 62%. Similarly, offer rates for pension accounts increase for all workers with incomes below the 6^{th} decile and deteriorate for those above. The lowest-income workers see their probability of finding a job offering a pension plan double, from 32% to 64%, while chances for workers in the top decile decline slightly, from 82% to 64%.

Figure 9 shows the resulting changes by income at the extensive and intensive margins for home equity and pension wealth compared to the baseline economy. Panel (a) reveals that participation in the housing market rises notably for all households up to the 6^{th} income decile, with the most significant increase occurring in the bottom decile, where participation jumps by a factor of 1.9 from 27% to 51%. In contrast, participation rates in the top three income deciles remain virtually unchanged. Although high-income households experience some deterioration in market access, this has only a negligible effect on their extensive margin. However, at

[26] In Appendix C.2, we document the results when removing frictions in the mortgage market while leaving frictions in the pension market unaffected. This experiment quantifies the extent of asset reallocation that occurs when one market becomes more accessible, for example, due to policy changes or innovations in financial markets.

the intensive margin in Panel (b), there is a downward shift in wealth-to-income ratios along the entire income distribution, with the most substantial reduction of 20% (from 2.5 to 2.0) in the bottom income decile. In the top decile, the intensive margin remains unchanged. The lower intensive margin at the bottom is a direct result of better market access: households that were largely excluded from the housing market in the baseline model purchase homes more frequently now, but on average, they buy smaller homes compared to the baseline.

Figure 9: Extensive and intensive margins in the home equity and pension markets in the baseline model and homogeneous access experiment.

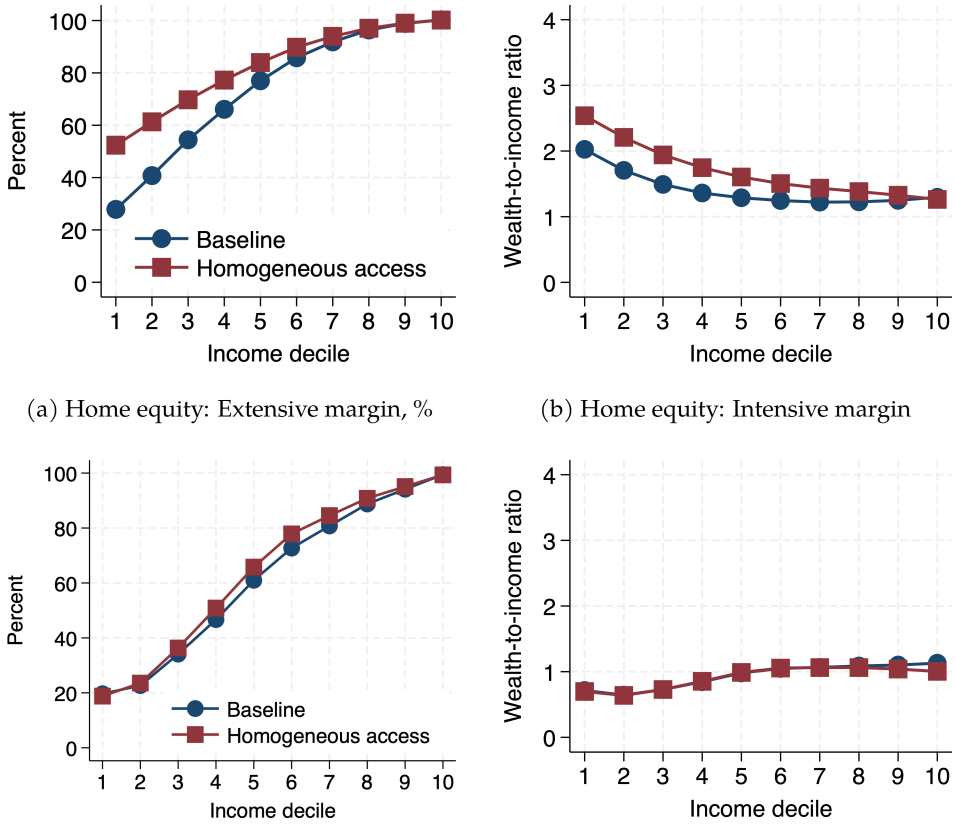

(a) Home equity: Extensive margin, %

(b) Home equity: Intensive margin

(c) Pension wealth: Extensive margin, %

(d) Pension wealth: Intensive margin

Notes: The homogeneous access experiment sets mortgage approval rates and offer rates for employer-sponsored retirement accounts to their economy-wide averages for all income levels (Figure 8).

Looking at the extensive margin for pension wealth in Figure 9(c), we find that the most significant changes happen in the middle of the income distribution, where participation increases, albeit only slightly. Note that we do not find any increase in participation in the pension market at the very bottom of the income distribution. The reason is the availability of the social security system that provides households with wealth in retirement and therefore lowers the incentives for life-cycle savings, consistent with Hubbard et al. (1995). Furthermore, we also find a modest increase in participation among households with above-median incomes, specifically those in the 6^{th} and 7^{th} income deciles. This increase results from better access to the pension market earlier in life when the agents' incomes are lower on average. In contrast, agents in the top two income deciles show no change in participation. Even though these agents face a roughly 10pp lower probability of being offered a retirement account, they receive a sufficiently large number of offers during the early and mid stages of their careers to get access to the pension market.

As for the intensive margin in the pension market, Figure 9(d) documents only a minor change. Recall that employer-sponsored pension plans in our framework are assumed to be pre-specified financial products, with each participant required to contribute the same percentage of their income. This assumption implies that the intensive margin of pension wealth should respond only insignificantly to changes in market access, primarily reflecting changes in the timing when agents start participating in the pension market rather than how much more they contribute.

Overall, we find that jointly improving asset market access leads to a stronger effect on the housing market. Moreover, since capital gains on housing are quite modest under our calibration, most of the wealth increase results from active savings, with less than 10% of the additional wealth accumulation accounted for by capital gains. In the next section, we analyze a policy experiment where only pension market access is expanded and only for lower-income households. This leads to a much stronger effect on wealth accumulation in the pension market, amplifying the role of capital gains in overall wealth accumulation.

Figure 10(a) shows the percentage change in wealth-to-income ratios across income deciles. We report the percentage change in wealth-to-income ratios for wealth, home equity, and pension wealth in the experiment relative to the respective ratio in the benchmark economy. Several observations emerge from the figure. First, there is a pronounced surge in home equity within the bottom half of the income distribution. This increase reflects the higher participation rates, as now more lower-income households buy homes, albeit they buy smaller ones on average. The total stock of home equity for this group increases, suggesting that the extensive margin effect (higher participation) dominates quantitatively over the intensive margin (smaller house size). Conversely, there is only a modest increase in home equity for agents in higher income brackets. Since participation remains virtually unchanged at the top, the observed decline results solely from these households purchasing smaller homes. These findings highlight the differential impact of homogeneous frictional access on home equity accumulation across the income distribution. Second, pension wealth weakly increases for everyone except those in the top two income deciles. This increase is due to higher participation rates at an approximately constant intensive margin. Combining the effects on home equity and pension wealth, wealth increases for households in the bottom half of the income distribution and decreases for those in the top income decile. As a result, wealth inequality between income groups declines.[27]

The decline of wealth inequality is clearly reflected in the Gini coefficient, which falls from 0.69 to 0.64. This 5-point decline is large at historical standards; for comparison, the Gini coefficient increased by only 4 points over the entire period from 1989 to 2022 in the SCF data (Kuhn and Ríos-Rull, 2025). How much the distribution of wealth changes in our experiment becomes also apparent when looking at the skewness of the distribution. The ratio of wealth in the top income decile to the bottom decile almost halves from 80 to 44. This shift is primarily driven by the bottom of the distribution moving closer to the middle. For example, the wealth ratio of the 5$^{\text{th}}$ to the 3$^{\text{rd}}$ decile declines from 2.7 to 1.5. At the same time, the middle also catches up somewhat with the top: the wealth ratio of the top decile to the 5$^{\text{th}}$ decile falls from 12 to

[27] Business wealth (not shown) rises, on average, by 8.5%.

Figure 10: Change in wealth-to-income ratios in the homogeneous access experiment.

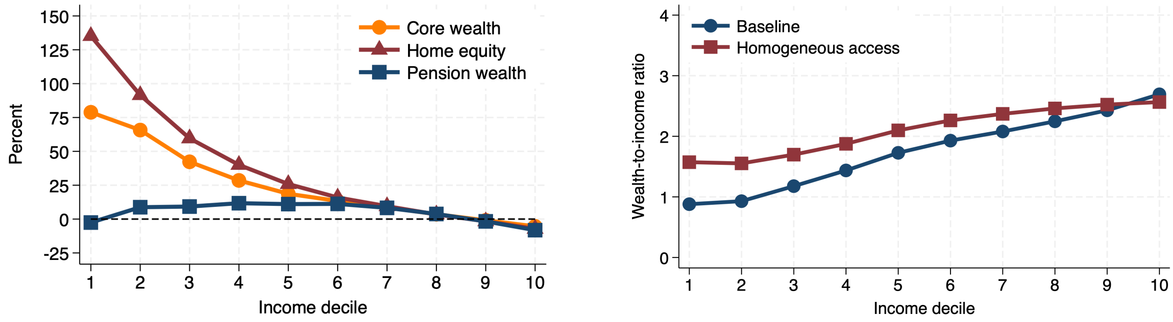

(a) Change in wealth-to-income ratios

(b) Wealth-to-income ratios

Notes: The homogeneous access experiment sets mortgage approval rates and offer rates for employer-sponsored retirement accounts to their economy-wide averages for all income levels. Left panel shows the percentage change in wealth-to-income ratios for wealth, home equity, and pension wealth in the homogeneous access experiment relative to the baseline. Right panel shows levels of the wealth-to-income ratios in the homogeneous access experiment and the baseline model.

9. Hence, we find a strong leveling effect from improving the wealth-building opportunities in the lower half of the income distribution.

This leveling effect can be clearly seen in Figure 10(b) that shows the level of wealth-to-income ratios in the homogeneous access experiment and the baseline model. Wealth-to-income ratios in the bottom half of the income distribution are 32% higher with homogeneous market access, while they are only 6% higher in the top half relative to the baseline model with heterogeneous access. The model with homogeneous access to asset markets leads to more wealth accumulation per dollar of income for lower-income households and thus a flatter profile of wealth-to-income ratios along the income distribution. This experiment demonstrates that the assumption of equal access to asset markets overstates wealth accumulation opportunities for low-income households and thereby understates wealth inequality across income groups.

Table 4: Saving rates by income in the homogeneous access experiment compared to the baseline model.

	Q1	Q2	Q3	Q4	Q5
Heterogeneous access (baseline)	0.008	0.085	0.141	0.159	0.173
Homogeneous access, pp.	+2.5	+4.4	+1.4	-0.7	-0.4

Notes: The table shows the estimation results of median regressions of saving rate on current income as in Dynan et al. (2004) (their Table 3, second column). The results are displayed in percentage point deviations from saving rates in the baseline model. See also notes to Table 3.

We have emphasized the importance of contractual savings flows for wealth accumulation in Section 2.4. The higher participation rates in our experiment imply that we should see more households with contractual saving flows in the lower part of the income distribution and thus with higher saving rates. Dynan et al. (2004) conjecture that this mechanism is an important driver of the empirically observed differences in saving rates. Our experiment provides evidence in favor of this conjecture. Table 4 reports the percentage point changes in saving rates by income quintiles. As a reference, we repeat the saving rates from the baseline model from

Table 3. Strikingly, we find that savings rates in the bottom quintile, which were virtually zero in the baseline model, increase by 2.5pp. We see the most pronounced increase in the saving rate of 4.4pp for the second income quintile, suggesting that many of these households had only limited access to the housing and pension markets despite their existing demand for these assets. By contrast, we find a negative change in the top quintile where the saving rate drops by 0.4pp. The lower saving rate for these households results from worse trading opportunities to upgrade their homes, which reduces saving rates in home equity at the intensive margin.

In summary, these results support the view that US asset markets provide an uneven playing field for wealth accumulation. The unequal opportunities due to heterogeneous market access manifest themselves in the data in low asset market participation and low wealth accumulation among low-income households. The lack of market access as a problem for the wealth accumulation of low-income households, highlighted by our analysis, is an important policy issue that is already actively discussed by policymakers. In the final step, we use our model to study the effects on wealth accumulation of one of the existing policy proposals to address the lack of access to employer-sponsored pension plans.

6 Policy experiment: Retirement Savings for Americans Act

There has been a longstanding debate in the macroeconomic literature on designing fully-funded pension systems that allow households to accumulate sufficient wealth for retirement. This literature has predominantly focused on policy measures targeted at household behavior to accumulate pension wealth, such as myopic behavior in financial planning, lack of financial literacy, liquidity constraints, and uninsurable idiosyncratic risks (Hubbard et al., 1995; Conesa and Krueger, 1999; Dynan et al., 2004; Nishiyama and Smetters, 2007; Ábrahám et al., 2024, 2025). By contrast, limited access to the market for retirement saving and the opportunity for households to build pension wealth has received no attention in the academic literature. The unequal access to employer-sponsored retirement accounts for lower-income workers is, however, an important public policy topic, and addressing this issue is at the heart of the recently proposed bipartisan *Retirement Savings for Americans Act*.[28] The Act proposes that workers without access to an employer-sponsored retirement plan become eligible for a newly introduced government-sponsored pension plan that provides income-dependent government copay for low-income workers. In this section, we will use our new model framework as a quantitative laboratory to study the consequences of this policy proposal for the wealth accumulation of income-poor workers.

6.1 Job and household variation in access rates

The results of any policy experiment conducted within a structural framework must be interpreted in light of the modeling assumptions. Therefore, before discussing these results, it is instructive to revisit two assumptions related to access and pickup rates for employer-sponsored retirement accounts.

First, our model assumes that access to employer-sponsored retirement plans depends on the

[28] A broad discussion of limited access for low-income households to the private pension market is provided by the *Economic Innovation Group* (EIG) at `https://inclusivewealth.eig.org/`.

worker's labor market income. The idea behind this modeling choice is that job-related differences are a key determinant of market access. In this sense, a job should be thought of as a bundle of the employer characteristics and other attributes such as the occupation. The assumption that access is job-related is important in the context of the policy experiment, as the government can then step in for employers who do not offer plans, for example, due to administrative costs.

To quantify the importance of the current job for access in the data, we statistically isolate the differences in income and access associated with the worker's job using our SCF dataset from the empirical section. We do so by regressing the variables of interest, income and access, on observable worker and job characteristics. Specifically, we run the following regression:

$$y_{i,t} = \alpha_t + \beta_W X_{W,t,i} + \beta_J X_{J,t,i} + \varepsilon_{i,t},$$

where $y_{i,t}$ denotes (log) income or the access indicator of household i in year t, α_t are year fixed effects, and $X_{W,t,i}$ and $X_{J,t,i}$ are worker and job observables of household i at time t. For job observables, we use dummies for 40 different industry-occupation combinations and five firm-size dummies. For worker characteristics, we use dummies to indicate if the household head is black, male, married, has a college education, and a third-order polynomial in age. We use hats to denote coefficient estimates and refer to $\hat{\beta}_J X_{J,t,i}$ as the job-related part of the outcome variable for household i and $\hat{\beta}_W X_{W,t,i}$ as the worker-related part. Then, we regress the job-related access part on the job-related income part. This yields a coefficient of 0.56, meaning a 10% increase in job-related income raises the access rate by 5.6pp. When we repeat this regression for the worker-related variables, the coefficient is 0.09. Thus, the job-related coefficient is nearly an order of magnitude larger than the worker-related coefficient.[29] We interpret this as evidence that job-related income differences are the primary factor in determining whether a worker receives an offer for a retirement plan or not. This evidence aligns at a more general level with a notion of "good" (high-paying) jobs (Acemoglu, 2001) that not only offer better pay but also offer better access to retirement accounts.

Second, regarding pickup rates, our assumption of stochastic access conditional on income implies that the pickup rate for workers who lack access to an employer-sponsored plan is the same as for those who have access. In other words, there is no selection based on observable worker characteristics. In the context of the policy experiment, a potential concern is that some employers may choose not to offer retirement plans if they expect their workers not to enroll in the offered plans (based on their observable characteristics). In this case, improving access for workers without access would be ineffective due to low pickup rates, which would compromise the results of our policy experiment. Although we consider endogenizing offer rates an important avenue for future research, we abstract from it here. Instead, we provide empirical support for our modeling approach by demonstrating that differences in observable worker characteristics cannot account for differences in (counterfactual) pick-up rates. The SCF only provides data on pickup for workers who have access to retirement plans. Since the counterfactual rates remain unobserved, we impute the counterfactual pick-up rates of workers without access based on their observable worker characteristics. If selection were based on observable worker char-

[29] In the raw data, the coefficient is 0.19. Detailed regression results are available from the authors upon request.

acteristics, the estimated pick-up rates would differ systematically between workers with and without access. However, we find that this is not the case, and that the predicted and observed pick-up rates align almost perfectly. We interpret this finding as strong support for high counterfactual pick-up rates for workers without access. For our policy experiment, this implies that we should expect a positive effect on participation if a government plan is introduced. However, the exact quantitative effect depends on the behavioral response of households and must be determined within the model, which is what we do next.

6.2 Design and implementation of the policy reform

To introduce the major provisions of the *Retirement Savings for Americans Act* into the model, we assume that workers who receive job offers without access to an employer plan can choose to enroll in a government plan.[30] Like private plans, the government plan is a predefined financial product. The contribution rate is fixed at $\chi_G = 3\%$ and worker income y determines the copay rate σ_G and a base rate $\bar{\chi}_G$, both financed by the government. For low-income workers who earn below half of the median income y_{med}, the government pays a base contribution rate of $\bar{\chi}_G = 1\%$ and matches the worker's contributions at rate $\sigma_G = 100\%$. Note that the reform requires workers to actively contribute to the account to qualify for the government base contribution. For workers with above-median income ($y > y_{med}$), the copay rate and the base rate are both set to zero. For agents with intermediate incomes $y \in [y_{med}/2, y_{med}]$, the base contribution rate $\bar{\chi}_G$ and the copay rate σ_G phase out linearly to zero. As for employer plans, the workers' contributions to government plans are tax-deductible.

Figure 11: Contribution rates to the pension accounts in the policy experiment.

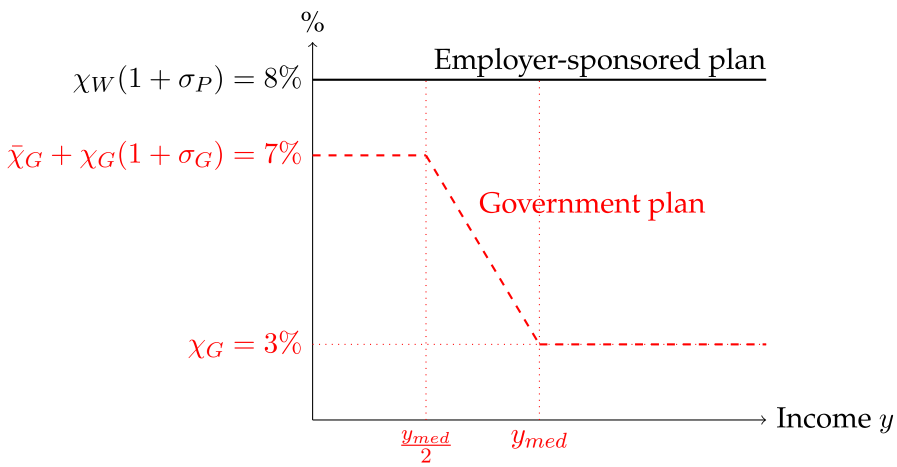

Notes: The solid line corresponds to the employer-sponsored plan, whose conditions are the same as in the baseline model (the worker's contribution is $\chi_W = 5.0\%$ and the employer copay is $\sigma_P = 59\%$). The dashed line corresponds to the government plan, where the worker's contribution is fixed at $\chi_G = 3\%$. The worker's labor income y determines the government copay rate σ_G and the base rate $\bar{\chi}_G$. For incomes below half of the median income $y_{med}/2$, $\bar{\chi}_G = 1\%$ and $\sigma_G = 1$. Both phase out linearly to 0% at the median income y_{med}.

Figure 11 illustrates the profile of contribution rates by income in the policy experiment. The dashed line corresponds to the government plan. The maximum contribution rate including copay for low-income workers is $\bar{\chi}_G + \chi_G(1+\sigma_G) = 7\%$. It phases out linearly from 7% to 3% for intermediate incomes and then remains constant at a contribution rate of 3% that is only paid

[30] Entrepreneurs are not eligible for government plans, and agents cannot enroll in both plans simultaneously.

by the worker. The solid line corresponds to the employer-sponsored plan, whose conditions are the same as in the baseline model with a cumulative contribution rate of $\chi_W(1+\sigma_P)$ that is constant at ca. 8% along the income distribution.

The top row of Figure 12 shows the pick-up rates for government and employer-sponsored plans in the model. The pick-up rate for government plans in Panel (a) is the percentage of workers with an active government plan among all eligible workers, which includes all workers as government plans are freely accessible. It is by construction zero in the baseline economy where this option is unavailable. The pick-up rate for employer plans in Panel (b) is the percentage of workers who have an active employer-sponsored plan among all workers eligible for the plan.

Figure 12: Pension market in the baseline model and policy experiment.

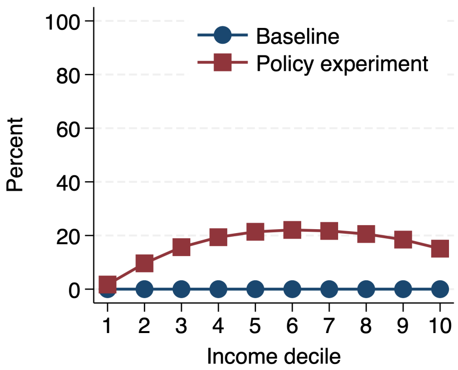

(a) Pick-up rates for government plans, %

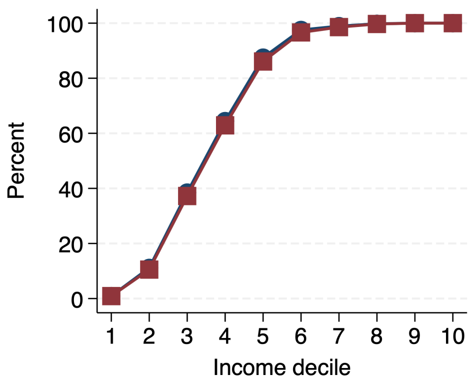

(b) Pick-up rates for employer-sponsored plans, %

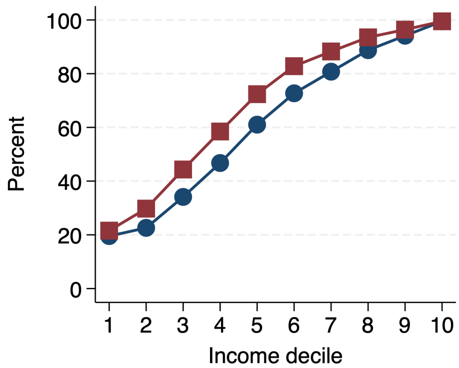

(c) Pension wealth: Extensive margin, %

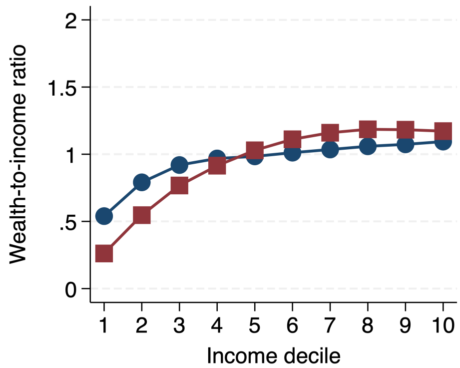

(d) Pension wealth: Intensive margin

Notes: In Panel (a), the pick-up rate is defined as the percentage of workers who hold an active government plan among all workers in the economy. In Panel (b), the pick-up rate is defined as the percentage of workers who hold an active employer-sponsored account among all workers eligible for an employer-sponsored plan.

Figure 12(a) shows that the policy boosts the demand for government plans along the entire income distribution. This effect comes entirely from the fact that the policy offers access to the pension market for all those workers who do not have access in the baseline model. Following the introduction of the policy, workers without access start entering into now-accessible government plans. Figure 12(b) reflects that the demand for employer-sponsored plans is not affected by the policy, indicating that there is no crowding out of private plans, the reason being that the option to choose between the government and the private plan does not to exist.

Figure 12(c) confirms that the policy reform significantly enhances participation in the pension market (+9pp). The most substantial increase is observed in the middle part of the income distribution, while participation remains virtually unchanged in the bottom income decile. The reason is again that the lowest-income households prefer not to save for retirement in the presence of social security (Hubbard et al., 1995).

Panel (d) of Figure 12 shows the intensive margin for pension wealth. We now see a reduction in the intensive margin for households with below-median incomes and an increase for higher-income households. Many households with below-median incomes enroll in government plans, and since these plans have lower contribution rates, pension wealth at the intensive margin declines on average. For high-income households, we find an increase in pension wealth at the intensive margin resulting from a life-cycle effect. When young, high-income households start in the lower parts of the income distribution and the newly available government-provided plan offers them earlier entry to the pension market so that when their income is high, they have already accumulated more pension wealth.

Figure 13 documents the effect of the policy on wealth accumulation by showing the percentage change in wealth-to-income ratios by income for total wealth, home equity, and pension wealth under the policy compared to the baseline model. Two important messages emerge from the figure. First, the policy has a pronounced impact on the pension wealth accumulation of households with below-median incomes. Looking at the entire bottom half of the income distribution, we see that pension wealth increases, on average, by 23%, suggesting that improving access to this market has a strong effect on their wealth accumulation. Yet, this percentage increase has to be interpreted with care as the level of retirement wealth for some low-income households has been very low prior to the reform. Second, we observe a weak substitution effect from the housing market to the pension market after access to the pension market has improved. The moderate decrease in housing wealth (on average, 2%) indicates that there is asset substitution, mainly occurring in the upper part of the income distribution where agents increase pension wealth due to the life-cycle effect, as we have just explained. Overall, the policy leads to an increase in wealth by 9%, with the most pronounced increase concentrated in the lower half of the income distribution.

Access to the pension market provides not only an opportunity for wealth accumulation but also the benefit of earning capital gains. Compared to the housing market, capital gains on pension wealth are more than five times larger under our calibration (Table 1). As a result, these additional capital gains amplify the wealth accumulation effect of improved savings opportunities. Specifically, we find that one in five dollars of the median household's additional wealth after the introduction of the government plan stems from capital gains.

Looking at the costs of implementing the reform, we find them to be fairly low, amounting to less than 1% of the economy-wide average pre-tax income. The largest government expense should fall on households with incomes below $y_{med}/2$ who qualify for both a full base rate $\bar{\chi}_G = 1\%$ and a full copay rate of $\sigma_G = 100\%$, but Figure 12(b) shows that pickup rates for government plans in the bottom income decile are less than 10%. Recall that the reform requires workers to actively contribute to the account to qualify for the government contribution, so the

Figure 13: Change in wealth-to-income ratios in the policy experiment compared to the baseline model, %.

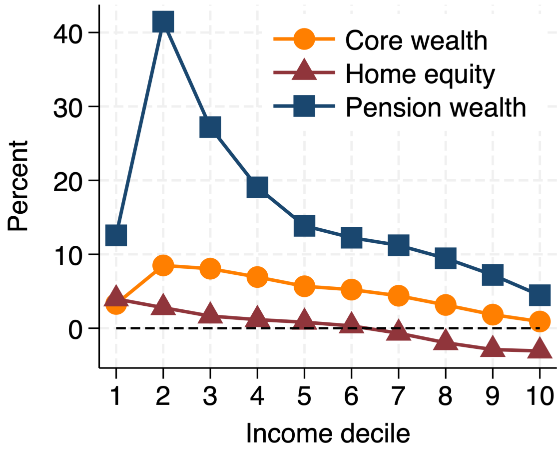

Notes: The figure shows the percentage change in wealth-to-income ratios for wealth, home equity, and pension wealth in the policy experiment relative to the respective ratio in the baseline model. Wealth is the sum of home equity, business wealth, and retirement accounts.

mandatory 3% contribution rate deters low-income agents from participating. These are lower middle-class households who actively participate in the government plans, but they qualify for a small subsidy only. The fact that it is mainly about access rather than subsidies keeps the macroeconomic costs limited.

Finally, we ask how policy affects the saving rates of low-income households. The comparative statics analysis from the previous section showed that saving rates of low-income households increase when market access improves. Table 5 confirms this finding in the context of the policy experiment. Throughout the income distribution, saving rates increase after the reform, consistent with the uniform increase in core wealth documented in Figure 13. The most pronounced increase in saving rates is observed in the bottom income quintile, where the savings rate rises by 1pp. The reason is that lower contribution rates in the government plan are particularly attractive for low-income households with weaker life-cycle savings motives. In the second to fourth income quintiles, savings rates rise on average by 76bp, while there is almost no change in the top income quintile due to the asset substitution effect described above.

Table 5: Saving rates by income in the policy experiment compared to the baseline model, pp.

	Q1	Q2	Q3	Q4	Q5
Policy Experiment	+1.01	+0.71	+0.78	+0.70	+0.01

Notes: The table shows the estimation results of median regressions of saving rate on current income as in Dynan et al. (2004) (their Table 3, second column). The results are displayed in percentage point deviations from saving rates in the baseline model. See also the note to Figure 3.

In summary, our experiment demonstrates that implementing a policy designed to alleviate limited access to the pension market significantly enhances broad-based wealth accumulation. By providing an alternative savings product with subsidies and lower contribution rates, the policy enables low-income households to build retirement savings in a government-provided plan. However, the bottom 10% of households remain largely unaffected by the policy, not because they do not have the opportunity to save, but because they do not want to save given the existing government programs in place.

6.3 External validity

The model suggests a sizable effect on wealth accumulation and household portfolios in response to a policy that improves access to the market for retirement plans. To put these quantitative results into perspective, we compare them to the evolution of employer-sponsored retirement plans during a comparable historical episode in the 1980s in the United States. This period marked a rapid growth in employer-sponsored retirement accounts, largely driven by the adoption of the 1981 *Economic Recovery Tax Act* (ERTA). Prior to the ERTA, retirement savings were limited to defined-benefit plans and Individual Retirement Accounts (IRAs). Moreover, strict eligibility requirements and limited flexibility in choosing plans restricted broader participation. ERTA was an instrumental policy that alleviated supply-side frictions in the pension market. First, it introduced 401(k) plans, which allowed employees to defer a portion of their income into retirement accounts on a tax-free basis, reducing their immediate tax burden. Second, the Act increased tax deductions available to employers for contributions, thereby making it financially more attractive for employers to offer such plans. Lastly, ERTA reduced administrative costs associated with setting up and managing pension plans. These provisions improved access to the pension market, very similar to our policy experiment discussed previously, so we take this historical policy as an opportunity to validate our model's predictions regarding the macroeconomic effects of such changes.

We rely on data from the SCF+ (Kuhn et al., 2020) and adhere to the same variable definitions and sample selection criteria as described in Section 2. In the SCF+ data, we verify that no wealth was reported for employer-sponsored plans before 1983.[31] Figure 14 depicts the dynamics of participation rates in the pension market in Panel (a) and the intensive margin for pension wealth in Panel (b) from 1983 to 1998 as blue squares across five survey waves. Following the ERTA adoption, participation in defined-contribution plans surged from 0% to 33%. The participation rates continued to rise thereafter, almost doubling to 64% until 1998. Meanwhile, the intensive margin for pension wealth also showed rapid growth, from 0.3 in 1983 to 0.7 in 1998. While participation rates for home equity remained virtually unchanged at 64% during this period (Panel (c)), the intensive margin for home equity declined by 28% from 1.8 to 1.3 (Panel (d)).

To assess whether our calibrated framework replicates the described trends in the historical data, we conduct the following stylized experiment. First, we solve the baseline model assuming that access to the pension market is entirely shut down by setting the probability of receiving an employer-sponsored retirement account, $\pi_P(y)$, to zero for all income levels (government plans are not offered in this experiment). We maintain all other parameters at their calibrated values. This situation represents the US economy in 1981, prior to the adoption of the ERTA. Then, we introduce an unanticipated and permanent reform, which sets the probabilities of accessing employer plans to the levels of the calibrated model. We simulate a panel of

[31] The SCF+ follows the SCF convention to report only marketable wealth, which excludes defined-benefit plans as they are not marketable. Note further that the SCF+ omits the 1986 survey.

Figure 14: Average extensive and intensive margins for retirement wealth and home equity in the simulation and data.

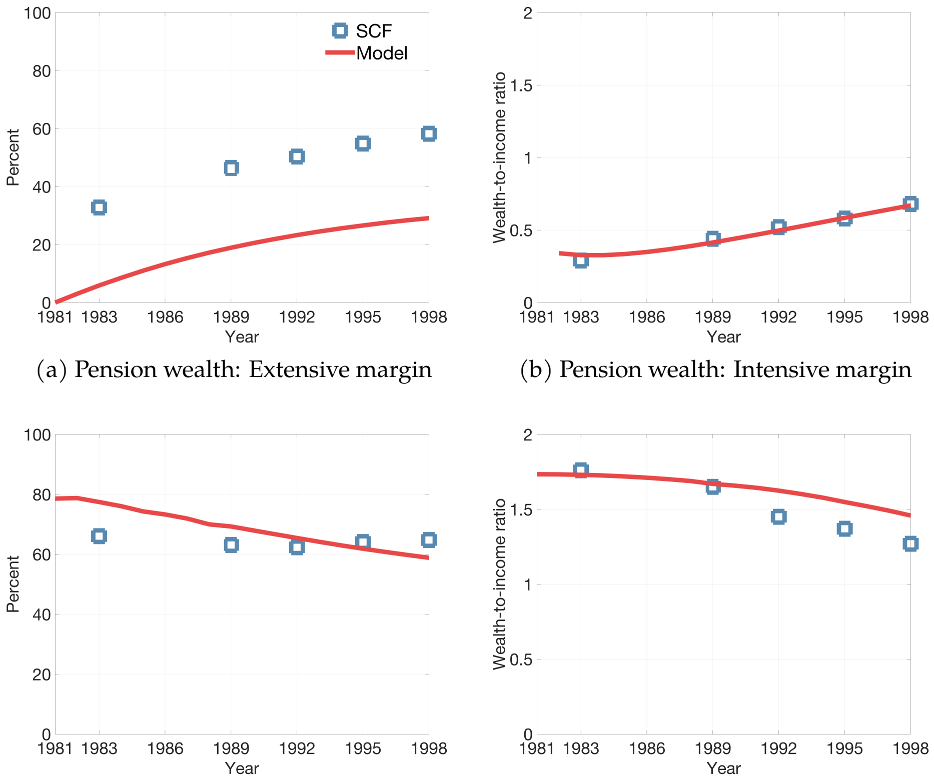

(a) Pension wealth: Extensive margin
(b) Pension wealth: Intensive margin
(c) Home equity: Extensive margin
(d) Home equity: Intensive margin

Notes: The figure shows the evolution of pension wealth and home equity in the US following the adoption of the 1981 Economic Recovery Tax Act. The empirical moments are taken from the SCF+ (Kuhn et al., 2020) and adhere to the same variable definitions and sample selection criteria as described in Section 2.

households applying the decision rules obtained from the solution of the baseline model and analyze the behavioral response of agents to changes in the supply of pension plans.

Figure 14 (solid line) shows the simulated paths for the extensive and intensive margins for pension wealth (Panels (a) and (b)) and home equity (Panels (c) and (d)). Reassuringly, we find that the model successfully captures the increasing participation rates and wealth-to-income ratios for pension wealth in the data. Quantitatively, the model explains 41% of the total growth in the extensive margin for pension wealth during 1983–1998. It mainly misses the strong initial increase after 1981 when probably more workers could directly opt into the new plans than in our stylized experiment. At the intensive margin, the model replicates remarkably closely the trajectory for the wealth-to-income ratios of pension wealth. For home equity, the model predicts a stable extensive margin and a decreasing wealth-to-income ratio, closely matching the empirical trend. The model accounts for 44% of the total drop in the intensive margin for home equity. The stronger decline at the extensive margin likely results from the simplification of our simulation, which assumes there are no pre-existing retirement plans that would lower wealth accumulation in the housing market. Overall, we take the fit of the model for this historical episode as additional evidence to support our quantitative findings.

7 Conclusions

This paper documents, both empirically and theoretically, the key role played by heterogeneity in access to asset markets in explaining differences in the accumulation of wealth among US households.

For the empirical analysis, we use the SCF data and focus on the three main asset classes in the portfolio of US households (home equity, business equity, and retirement accounts). Using these data, we provide three new facts about the distribution of wealth. First, we document that differences in the extensive margin along the income distribution are key in accounting for the unequal distribution of wealth. Second, we provide evidence that this heterogeneity is closely related to variation in asset market access. Third, we document the key role of contractual savings for wealth building of US households.

Guided by these new empirical facts, we develop a new model of life-cycle wealth accumulation and lifetime financial decision. The model's wealth accumulation mechanism focuses on the participation in asset markets (extensive margin) rather than on the variation in the savings of asset market participants (intensive margin). Households in the model execute pre-specified financial contracts and have consumption dynamics like wealthy hand-to-mouth households, which are typically not on their Euler equation.

We show that this model accounts for the empirical evidence on the joint distribution of income and wealth and our new evidence on asset market participation. Greater wealth accumulation by higher-income households in the model does not result from higher shares of income saved by households participating in asset markets but from differences in asset market participation. Capital gains of asset market participants amplify these differences. Our counterfactual analysis demonstrates that low-income households are typically constrained in their wealth accumulation due to a lack of access to asset markets. After market access improves, they begin to accumulate more wealth, and wealth differences by income decline. The conclusion is that US asset markets do not provide a level playing field for all households, as opportunities to build wealth vary systematically with income.

Finally, we conduct a policy experiment that analyzes the consequences of the *Retirement Savings for Americans Act* for wealth accumulation. The Act aims to improve market access for low-income households in the pension market. We find that the targeted low-income households accumulate substantially more retirement wealth after their market access improves. We validate the model predictions from the reform to a similar historical episode of the proliferation of employer-sponsored retirement accounts in the 1980s. The policy experiment corroborates the result of the important role of income-dependent market access for the wealth-building of American households.

References

Ábrahám, A., P. Brendler, and E. Carceles-Poveda (2024). Optimal income redistribution. *Working Paper*.

Ábrahám, A., P. Brendler, and E. Carceles-Poveda (2025). On taxation of private retirement wealth. *Working Paper*.

Acemoglu, D. (2001). Good jobs versus bad jobs. *Journal of Labor Economics 19*(1), 1–21.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics 109*(3), 659–684.

Alan, S., K. Atalay, and T. F. Crossley (2015). Do the rich save more? Evidence from Canada. *Review of Income and Wealth 61*(4), 739–758.

Bach, L., L. E. Calvet, and P. Sodini (2020, September). Rich pickings? Risk, return, and skill in household wealth. *American Economic Review 110*(9), 2703–47.

Bayer, C., B. Born, and R. Luetticke (2024). Shocks, frictions, and inequality in U.S. business cycles. *American Economic Review 114*(5), 1211–1247.

Bayer, C., R. Lütticke, L. Pham-Dao, and V. Tjaden (2019). Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk. *Econometrica 87*(1), 255–290.

Benhabib, J., A. Bisin, and M. Luo (2019). Wealth distribution and social mobility in the U.S.: A quantitative approach. *The American Economic Review 109*(5), 1623–1647.

Bewley, T. (1986). Stationary monetary equilibrium with a continuum of independently fluctuating consumers. *Contributions to mathematical economics in honor of Gérard Debreu 79*.

Bhattacharya, N. and B. Mahalanobis (1967). Regional disparities in household consumption in India. *Journal of the American Statistical Association 62*(317), 143–161.

Bhutta, N., J. Bricker, A. C. Chang, L. J. Dettling, S. Goodman, J. W. Hsu, K. B. Moore, S. Reber, A. Henriques Volz, and R. Windle (2020). Changes in US family finances from 2016 to 2019: Evidence from the Survey of Consumer Finances. *Federal Reserve Bulletin 106*(5).

Bilbiie, F. O. (2024). Monetary policy and heterogeneity: An analytical framework. *The Review of Economic Studies 92*(4), 2398–2436.

Bloom, N., C. I. Jones, J. Van Reenen, and M. Webb (2020). Are ideas getting harder to find? *American Economic Review 110*(4), 1104–44.

Brendler, P. (2023). Rising earnings inequality and optimal income tax and social security policies. *Journal of Monetary Economics 134*, 35–52.

Cagetti, M. (2003). Wealth accumulation over the life cycle and precautionary savings. *Journal of Business & Economic Statistics 21*(3), 339–353.

Cagetti, M. and M. De Nardi (2006). Entrepreneurship, frictions, and wealth. *Journal of Political Economy 114*(5), 835–870.

Cagetti, M. and M. De Nardi (2008). Wealth inequality: Data and models. *Macroeconomic Dynamics 12*(S2), 285–313.

Castañeda, A., J. Díaz-Giménez, and J. Ríos-Rull (2003). Accounting for the U.S. earnings and wealth inequality. *Journal of Political Economy 111*(4), 818–857.

Choukhmane, T. (2025). Default options and retirement saving dynamics. *American Economic Review* (*Conditionally accepted*).

Choukhmane, T., J. Colmenares, C. O'Dea, J. L. Rothbaum, and L. D. Schmidt (2024). Who benefits from retirement saving incentives in the U.S.? Evidence on gaps in retirement wealth accumulation by race and parental income. Working Paper 32843, NBER.

Choukhmane, T. and T. de Silva (2024). What drives investors' portfolio choices? Separating risk preferences from frictions. Working Paper 32476, NBER.

Conesa, J. C. and D. Krueger (1999). Social Security Reform with Heterogeneous Agents. *Review of Economic Dynamics 2*(4), 757–795.

De Nardi, M. (2004). Wealth inequality and intergenerational links. *Review of Economic Studies 71*(3), 743–768.

De Nardi, M. and G. Fella (2017). Saving and wealth inequality. *Review of Economic Dynamics 26*, 280–300.

De Nardi, M., E. French, and J. B. Jones (2010). Why do the elderly save? The role of medical expenses. *Journal of Political Economy 118*(1), 39–75.

Dynan, K. E., J. Skinner, and S. P. Zeldes (2004). Do the rich save more? *Journal of Political Economy 112*(2), 397–444.

Fagereng, A., M. B. Holm, B. Moll, and G. Natvik (2021). Saving behavior across the wealth distribution: The importance of capital gains. *Working Paper*.

Gomes, F., M. Haliassos, and T. Ramadorai (2021). Household finance. *Journal of Economic Literature 59*(3), 919–1000.

Gourinchas, P.-O. and J. A. Parker (2002). Consumption over the life cycle. *Econometrica 70*(1), 47–89.

Guiso, L., M. Haliassos, and T. Jappelli (2002). *Household portfolios*. MIT press.

Guiso, L. and P. Sodini (2013). Chapter 21 - Household finance: An emerging field. Volume 2 of *Handbook of the Economics of Finance*, pp. 1397–1532. Elsevier.

Haliassos, M. and C. C. Bertaut (1995). Why do so few hold stocks? *Economic Journal 105*(432), 1110–1129.

Heathcote, J., K. Storesletten, and G. L. Violante (2020). Presidential address 2019: How should tax progressivity respond to rising income inequality? *Journal of the European Economic Association 18*(6), 2715–2754.

Heikkuri, V.-M. and M. Schief (2025). Subgroup decomposition of the gini coefficient: A new solution to an old problem. *Econometrica* (*conditionally accepted*).

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1995). Precautionary saving and social insurance. *Journal of Political Economy 103*(2), 360–399.

Hubmer, J., P. Krusell, and A. A. Smith. (2021). Sources of US wealth inequality: Past, present, and future. *NBER Macroeconomics Annual 35*, 391–455.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control 17*(5-6), 953–969.

Jordà, Ò., K. Knoll, D. Kuvshinov, M. Schularick, and A. M. Taylor (2019). The rate of return on everything, 1870–2015. *The Quarterly Journal of Economics 134*(3), 1225–1298.

Jung, P. and M. Kuhn (2019). Earnings losses and labor mobility over the life cycle. *Journal of the European Economic Association 17*(3), 678–724.

Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary policy according to HANK. *American Economic Review 108*(3), 697–743.

Kaplan, G. and G. L. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica 82*(4), 1199–1239.

Kaplan, G., G. L. Violante, and J. Weidner (2014). The wealthy hand-to-mouth. *V Brookings Papers on Economic Activity, Spring*.

Krueger, D., K. Mitman, and F. Perri (2016). Macroeconomics and household heterogeneity. In *Handbook of macroeconomics*, Volume 2, pp. 843–921. Elsevier.

Krusell, P. and A. A. Smith, Jr (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy 106*(5), 867–896.

Kuhn, M. and J.-V. Rıos-Rull (2016). 2013 update on the U.S. earnings, income, and wealth distributional facts: A view from macroeconomics. *Federal Reserve Bank of Minneapolis Quarterly Review 37*(1).

Kuhn, M. and J.-V. Ríos-Rull (2025). Income and wealth inequality in the United States: An update including the 2022 wave. Technical report, National Bureau of Economic Research.

Kuhn, M., M. Schularick, and U. I. Steins (2020). Income and wealth inequality in America, 1949–2016. *Journal of Political Economy 128*(9), 3469–3519.

Low, H., C. Meghir, and L. Pistaferri (2010). Wage risk and employment risk over the life cycle. *American Economic Review 100*(4), 1432–1467.

Luetticke, R. (2021). Transmission of monetary policy with heterogeneity in household portfolios. *American Economic Journal: Macroeconomics 13*(2), 1–25.

Mankiw, N. G. and S. P. Zeldes (1991). The consumption of stockholders and nonstockholders. *Journal of Financial Economics 29*(1), 97–112.

Nishiyama, S. and K. Smetters (2007). Does Social Security privatization produce efficiency gains? *The Quarterly Journal of Economics 122*(4), 1677–1719.

Quadrini, V. (2000). Entrepreneurship, saving, and social mobility. *Review of Economic Dynamics 3*(1), 1–40.

Saez, E. and G. Zucman (2014). Wealth inequality in the United States since 1913: Evidence from capitalized income tax data. NBER Working Papers 20625, National Bureau of Economic Research, Inc.

Straub, L. (2019). Consumption, savings, and the distribution of permanent income. *Unpublished manuscript, Harvard University 17*.

Online Appendix

A Empirical analysis

A.1 Variable definitions and sample selection criteria

We use the publicly available data from the Survey of Consumer Finances (SCF). Our main analysis spans the years 2010 to 2019, but we also conduct a robustness check pooling all publicly available datasets since 1989 below. For the results on the relationship between income and wealth, we rely on the Summary Extract file. When analyzing access to the pension and home equity markets as well as the role of prices for asset-market participation, we resort to the full dataset with all variables.

A small set of sample selection restrictions is applied. First, we focus on household heads aged 25–65. Second, we restrict the sample to those households whose head works in dependent employment or is self-employed. Third, we drop households that report running a business but not actively managing it.

Business wealth is the sum of businesses in which the household has an active interest, and it is computed as the value of the business if sold net of all liabilities. Home equity comprises both primary residence and other residential real estate (e.g., vacation homes) minus housing debts, such as mortgages and home equity lines of credit. Retirement accounts include pension plans accumulated in individual retirement and thrift accounts, net of any loans taken against pensions. To cover the remainder of the household balance sheet, we construct the following three asset classes: (i) stock, bonds, and mutual funds, (ii) liquid assets, and (iii) other assets. Liquid assets comprise all types of transaction accounts (checking, savings, money market accounts, etc.) plus certificates of deposit net of credit lines, credit balances, etc. Other assets include remaining financial and real assets such as the cash value of life insurance, vehicles, etc., net of any corresponding debts.

We compute wealth-to-income ratios for each household and report the average ratio within each income decile. But our results remain very similar if we compute the ratio of average wealth to average income by decile. To exclude outliers, we winsorize the top 1% of the wealth-to-income ratios for total wealth and each core wealth component and then construct wealth-to-income ratios for core wealth from the winsorized data.

When analyzing access to the pension market and the price effects for participation in the pension market, we impose an additional sample selection restriction that the household head works in dependent employment. However, income deciles are defined on a full sample to maintain consistency with the preceding analysis.

A.2 Alternative sample period

Figure 15 confirms that the documented patterns for total wealth and core wealth, shown in Figure 1 in the main text for 2010–2019, also hold when we pool all publicly available SCF datasets since 1989. Similarly, Figure 16 shows that the documented patterns for the extensive and intensive margins (Figure 2) are also robust with respect to the sample period.

Figure 15: Wealth-to-income ratios for total wealth and core wealth in the SCF (1989–2019).

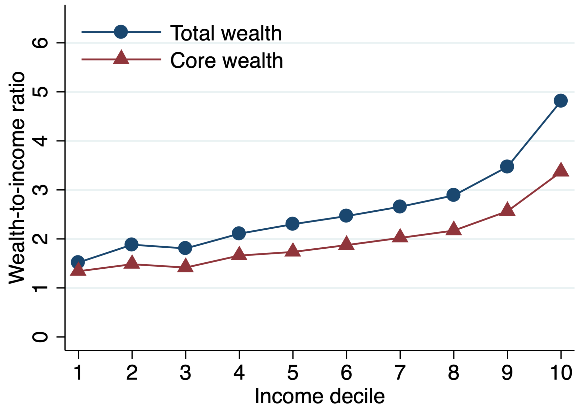

Notes: The figure is constructed using the full publicly available SCF dataset covering 1989–2010. *Core wealth* is the sum of business wealth, home equity, and retirement accounts. *Total wealth* comprises core wealth plus stocks, bonds, mutual funds, liquid assets, and remaining assets. The value of each asset component is computed net of debts. See Appendix A.1 for the sample selection criteria.

Figure 16: Extensive margin and intensive margin variations by income in the SCF (1989–2019).

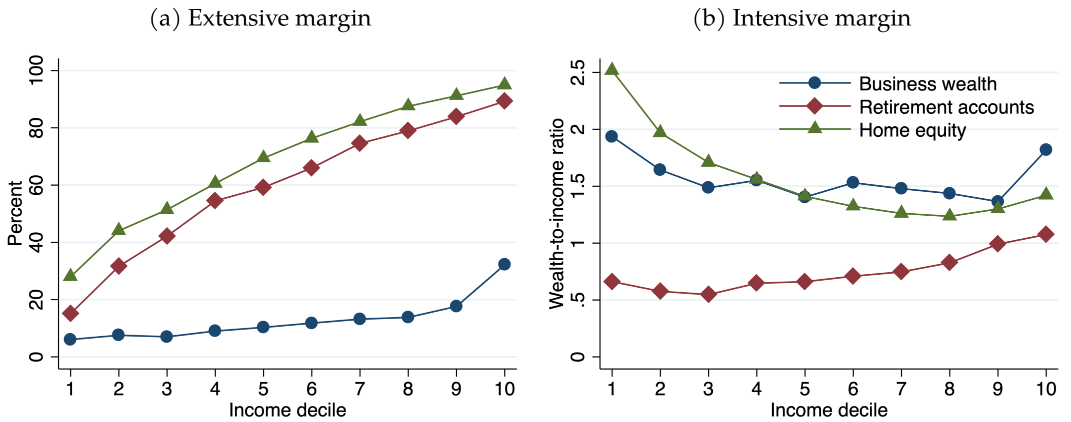

Notes: The figure is constructed using the full publicly available SCF dataset covering 1989–2010. See Appendix A.1 for the sample selection criteria.

A.3 Portfolio composition

Table 6 reports the portfolio shares in total wealth held by three income groups (0–50%, 50–90%, and 90–100% of the income distribution). In each column, the portfolio shares sum up to 100% for a given income group.

The last column shows the economy-wide average shares of each asset class in total wealth. The portfolio components with the largest shares are business wealth (30.6%), home equity (21.4%), and retirement accounts (17.7%). Together, these three asset classes constitute core wealth and comprise 70% of household wealth. The remaining wealth components and their shares in total wealth are: Stocks, bonds, and mutual funds (15.6%), Liquid assets (6.1%), and Other assets (8.6%). Stocks, bonds, and mutual funds comprise the largest asset class, not

Table 6: Portfolio composition by income groups in the SCF, %.

	Wealth component	Share of asset held by income group:			
		$< 50\%$	$50-90\%$	$> 90\%$	All
	Business wealth	18.5	16.1	33.9	30.6
+	Home equity	41.2	31.5	18.5	21.4
+	Retirement accounts	17.2	28.2	15.9	17.7
=	**Core wealth**	**77.0**	**75.8**	**68.3**	**69.8**
+	Stocks, bonds, and mutual funds	8.8	8.6	17.3	15.6
+	Liquid assets	5.6	6.3	6.0	6.1
+	Other assets	8.7	9.3	8.4	8.6
=	**Total wealth**	**100.0**	**100.0**	**100.0**	**100.0**

included in core wealth, but these assets are heavily concentrated at the top of the distribution.[32]

Looking at the portfolio shares along the income distribution displayed in the three middle columns, we see that for each income group, the sum of business wealth, home equity, and pension wealth accounts for the largest share of their total wealth. For the bottom 90% of households, the share of these three portfolio components is even higher than the average and virtually the same for the first two income groups covering the bottom 90% (77.0% and 75.8%). For the top 10%, the share is only slightly below 70%. As one can see from the last rows of the table, the key difference between the top 10% and the bottom 90% is the higher share of stocks, bonds, and mutual funds, which is below 9% for the bottom 90% of the income distribution and almost twice as high (17.3%) for the top 10%. The largest share of Other assets are cars, and together with liquid assets, which primarily consist of transaction accounts, they account for nearly two-thirds of the non-core wealth of the bottom 90%. Note that these low/no-return asset classes are typically not used to accumulate life-cycle savings.

A.4 Non-core wealth

In this section, we report wealth-to-income ratios for non-core wealth by income and over the life cycle.

The left panel of Figure 17 plots wealth-to-income ratios for non-core wealth by income and demonstrates that non-core wealth is substantially lower and hardly varies between the 2^{nd} and 9^{th} income deciles. Only the top decile has substantial holdings of non-core wealth consistent with stock holdings concentrated in that group. The right panel of Figure 17 further dissects the wealth-to-income ratios for non-core wealth presented in the left panel by categorizing households into two groups: those with positive amounts of core wealth and those without any core wealth. We can see that households lacking core wealth also hold virtually no non-core wealth. Put together, our findings indicate that there is no evidence to suggest a substitution effect from core to non-core wealth among households without core wealth.

Figure 18 plots wealth-to-income ratios for total wealth, core wealth, and non-core wealth over

[32] Kuhn et al. (2020) document large and persistent differences in portfolio composition of US households along the wealth distribution and report that the top 10% of households own more than 90% of all stocks.

Figure 17: Joint distribution of non-core wealth and income in the SCF.

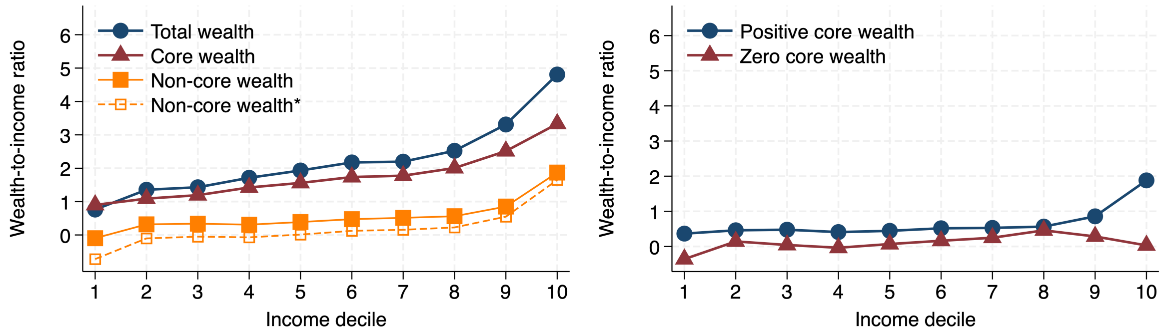

(a) Core and non-core wealth. (b) Non-core wealth by core wealth holdings.

Notes: The left panel shows wealth-to-income ratios for total wealth, core wealth, and non-core wealth by income. The right panel shows wealth-to-income ratios for non-core wealth by income for households with positive holdings of core wealth and for households with zero holdings of core wealth. The figure is constructed using the full publicly available SCF dataset covering 1989–2010. *Core wealth* is the sum of business wealth, home equity, and retirement accounts. *Total wealth* comprises core wealth plus stocks, bonds, mutual funds, liquid assets, and remaining assets. *Non-core wealth* is Total wealth net of Core wealth. *Non-core wealth** is Non-core wealth net of vehicles and checking accounts. The value of each asset component is computed net of debts. See Appendix A.1 for the sample selection criteria.

the life cycle. The figure shows that roughly 80% of the accumulated wealth at the end of the life cycle is core wealth. Over the life cycle, core wealth-to-income ratios increase from about 0.5 to 4, whereas non-core wealth increases from zero to one. Hence, the largest share of life-cycle wealth accumulation is accounted for by core wealth. This finding is consistent with our model in Section 3, which focuses on life-cycle wealth accumulation.

Figure 18: Non-core wealth over the life cycle in the SCF.

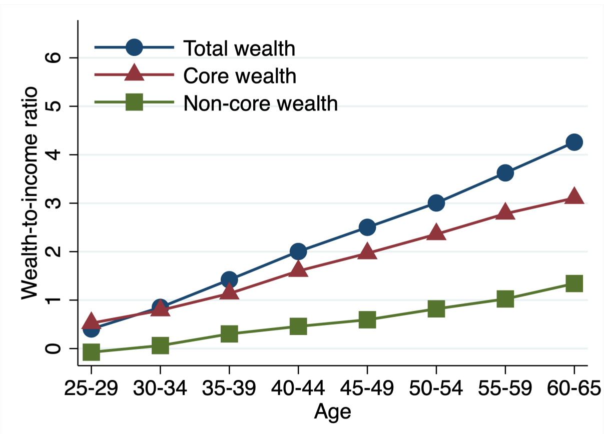

Notes: The figure is constructed using the full publicly available SCF dataset covering 1989–2010. *Core wealth* is the sum of business wealth, home equity, and retirement accounts. *Total wealth* comprises core wealth plus stocks, bonds, mutual funds, liquid assets, and remaining assets. *Non-core wealth* is Total wealth net of Core wealth. The value of each asset component is computed net of debts. See Appendix A.1 for the sample selection criteria.

A.5 Sensitivity analysis

In this subsection, we demonstrate that the relationship between income and wealth, documented in Figure 1, remains economically and statistically significant, even when accounting for a broad set of household characteristics. Additionally, we perform sensitivity analyses on the income dependence of mortgage approval rates and offer rates for employer-sponsored retirement plans to show that conditional on various household characteristics, income deciles exert a statistically significant and positive effect on both variables.

Section A.5.1 presents the regression results based on our primary SCF dataset. To corroborate our results on income dependence of market access, Section A.5.3 uses *Home Mortgage Disclosure Act* (HMDA) data to study the relationship between mortgage approval rates and income. Section A.5.4 relies on data from the *National Compensation Survey* (NCS) to corroborate the evidence for employer-provided pension plans. Finally, Section A.5.2 addresses some concerns regarding the role of prices for the extensive margin for pension wealth and home equity.

A.5.1 Survey of Consumer Finances

Variable definitions

Below we define the main variables used in our analysis (see the SCF definitions for other variables):

- *Offered* – An indicator variable that equals one if a household head has access to an employer-provided pension plan and zero otherwise. A household head is defined to have access if the following two conditions hold: their employer offers such a pension plan (*X4136*) and they are eligible to enroll in the plan (*X4137*).
- *Applied* – An indicator variable that equals one if at least one of the following three conditions holds: the household applied for a mortgage or a home-based loan (*X435*), or requested to refinance a mortgage (*X436*), or did not apply for any credit because it did not expect to get approved (*X441*). The indicator variable equals zero when the household responds *No* to all three questions.
- *Rejected* – An indicator variable that equals one if variable *Applied* equals one and at least one of the following two conditions is fulfilled: the request for credit is turned down, the household received less credit than what was applied for (*X407*), or the household did not apply for any credit because it did not expect to get approved (*X441*). The indicator variable equals zero when variable *Applied* equals one and no credit application was rejected in the last twelve months (*X407*).
- *Approved* – An indicator variable that equals one if the variable *Rejected* equals zero and zero if *Rejected* equals one.
- *Income Decile Dummies* – A set of indicator variables that indicate the household head's income decile. The first income decile constitutes the reference group for the probit regressions. In all figures that rely on the SCF data, the boundaries of the income deciles are kept the same throughout the paper to preserve consistency.
- *Age* – A variable between 25 and 65 that indicates the household head's age. When included as a control variable, we always use both age and age squared.

- *College* – An indicator variable that equals one if the household head holds at least some college or associate degree and zero otherwise. We use the categorical variable *edcl* to determine the household head's educational level.

Results

Table 7 presents the OLS regression results of wealth-to-income ratios on a broad set of household characteristics. The table confirms that the rising pattern of wealth-to-income ratios for total and core wealth documented in Figure 1 in the main text is not fully accounted for by the household life-cycle saving motives and that the household income has a statistically significant effect on the wealth-to-income ratios, conditional on other observable characteristics.

Table 7: OLS regressions of wealth-to-income ratios (SCF).

	(1) Total	(2) Total	(3) Core	(4) Core
Income decile:				
-2nd	0.596	0.531	0.188	0.131
	(0.183)	(0.171)	(0.095)	(0.084)
-3rd	0.672	0.566	0.294	0.202
	(0.182)	(0.176)	(0.083)	(0.076)
-4th	0.947	0.696	0.517	0.316
	(0.161)	(0.154)	(0.075)	(0.070)
-5th	1.161	0.846	0.650	0.399
	(0.180)	(0.175)	(0.079)	(0.066)
-6th	1.405	0.961	0.827	0.483
	(0.147)	(0.146)	(0.065)	(0.061)
-7th	1.416	0.931	0.859	0.471
	(0.169)	(0.168)	(0.071)	(0.067)
-8th	1.753	1.192	1.097	0.648
	(0.156)	(0.159)	(0.071)	(0.077)
-9th	2.545	1.811	1.607	1.026
	(0.158)	(0.160)	(0.083)	(0.082)
-10th	4.011	2.947	2.394	1.566
	(0.217)	(0.208)	(0.110)	(0.104)
Constant	0.756	-1.434	0.897	-1.046
	(0.151)	(0.473)	(0.060)	(0.200)
Demographic controls		x		x
Year controls		x		x
Observations	62,290	62,290	62,859	62,859

Notes: The regressions are based on the SCF 2010–2019 datasets. Demographic controls include age, age squared, gender, race, education, and marital status. Standard errors are reported in brackets below each coefficient estimate. See Appendix A.1 for the sample selection criteria.

Next, we perform sensitivity analyses on the income dependence of the extensive margin and market access. Table 8 shows the results of the probit regressions of the extensive margin in each asset component of core wealth on income deciles and a broad set of control variables. Table 9 presents the results of the probit regressions of mortgage approval rates, while Table 10 displays the results of the probit regressions of access to employer-provided pension plans.

The tables confirm that conditional on various household characteristics, income deciles exert a statistically significant and positive effect on the extensive margin and access rates.

Table 8: Probit regression of participation rates in the pension, home equity, and business equity markets (SCF).

	(1) Pensions	(2) Home equity	(3) Business equity
Income decile:			
-2nd	0.507	0.344	0.070
	(0.048)	(0.061)	(0.060)
-3rd	0.908	0.587	-0.005
	(0.049)	(0.054)	(0.078)
-4th	1.192	0.746	0.081
	(0.058)	(0.056)	(0.061)
-5th	1.294	0.993	0.152
	(0.043)	(0.050)	(0.070)
-6th	1.509	1.126	0.198
	(0.056)	(0.049)	(0.065)
-7th	1.820	1.376	0.204
	(0.050)	(0.058)	(0.067)
-8th	1.963	1.474	0.193
	(0.054)	(0.051)	(0.067)
-9th	2.170	1.713	0.362
	(0.061)	(0.070)	(0.065)
-10th	2.465	1.997	0.694
	(0.061)	(0.063)	(0.061)
Constant	-0.972	-2.284	-3.094
	(0.193)	(0.144)	(0.258)
Demographic controls	x	x	x
Year controls	x	x	x
Observations	66,345	66,454	67,995

Notes: The regressions are based on the SCF 2010–2019 datasets. Demographic controls include age, age squared, gender, race, education, and marital status. Standard errors are reported in brackets below each coefficient estimate. See Appendix A.1 for the sample selection criteria.

Table 9: Probit regression of mortgage approval rates (SCF).

	(1) Approved	(2) Approved
Income decile:		
-2nd	0.756	0.685
	(0.171)	(0.164)
-3rd	0.825	0.787
	(0.179)	(0.179)
-4th	1.142	1.058
	(0.164)	(0.158)
-5th	1.393	1.245
	(0.153)	(0.166)
-6th	1.589	1.430
	(0.180)	(0.183)
-7th	1.996	1.797
	(0.201)	(0.206)
-8th	1.825	1.553
	(0.147)	(0.160)
-9th	2.052	1.791
	(0.177)	(0.193)
-10th	2.196	1.876
	(0.202)	(0.210)
Constant	-0.994	-0.071
	(0.127)	(0.564)
Demographic controls		x
Year controls		x
Observations	7,087	7,087

Notes: The regressions are based on the SCF 2010–2019 datasets. Standard errors are reported in brackets below each coefficient estimate. An applicant is a household who submitted a loan application in the past 12 months or who did not submit an application in the past 12 months because he/she thought it would be denied. An application is treated as approved if the household receives full funding. An application is treated as rejected if the request for credit is turned down, the household receives less credit than what was applied for, or the household did not apply for any credit because it thought it would be denied. Applications for a mortgage, home-based loan, and refinancing are considered. Demographic controls include age, age squared, gender, race, education, and marital status. See the variable definitions above for more details. See the variable definitions in Appendix A.5.1. See Appendix A.1 for the sample selection criteria.

Table 10: Probit regression of offer rates for employer-sponsored retirement accounts (SCF).

	(1) Offered	(2) Offered
Income decile:		
-2nd	0.512	0.545
	(0.099)	(0.098)
-3rd	0.828	0.893
	(0.093)	(0.094)
-4th	1.041	1.141
	(0.093)	(0.094)
-5th	1.309	1.447
	(0.086)	(0.088)
-6th	1.436	1.598
	(0.089)	(0.093)
-7th	1.501	1.694
	(0.094)	(0.101)
-8th	1.666	1.879
	(0.086)	(0.092)
-9th	1.796	2.000
	(0.082)	(0.086)
-10th	1.871	2.057
	(0.093)	(0.100)
Constant	-0.846	-0.220
	(0.077)	(0.247)
Demographic controls		x
Year controls		x
Observations	37,715	37,715

Notes: The regressions are based on the SCF 2010–2019 datasets. Standard errors are reported in brackets below each coefficient estimate. A household is considered to be offered an employer-sponsored plan if the household head works for an employer who offers a retirement plan and is eligible to participate in the plan. Demographic controls include age, age squared, gender, race, education, and marital status. See the variable definitions in Appendix A.5.1. See Appendix A.1 for the sample selection criteria.

A.5.2 Contribution rates, copay rates, and mortgage interest rates

One reason for heterogeneity in participation rates in employer-provided retirement plans, documented in Figure 2(a) in the main text, could be differences in copay rates offered by employers. We rely on data from the SCF to study copay rates by employers for active plans along the income distribution.[33]

Figure 19: Copay and contribution rates for active employer-provided retirement plans in the SCF.

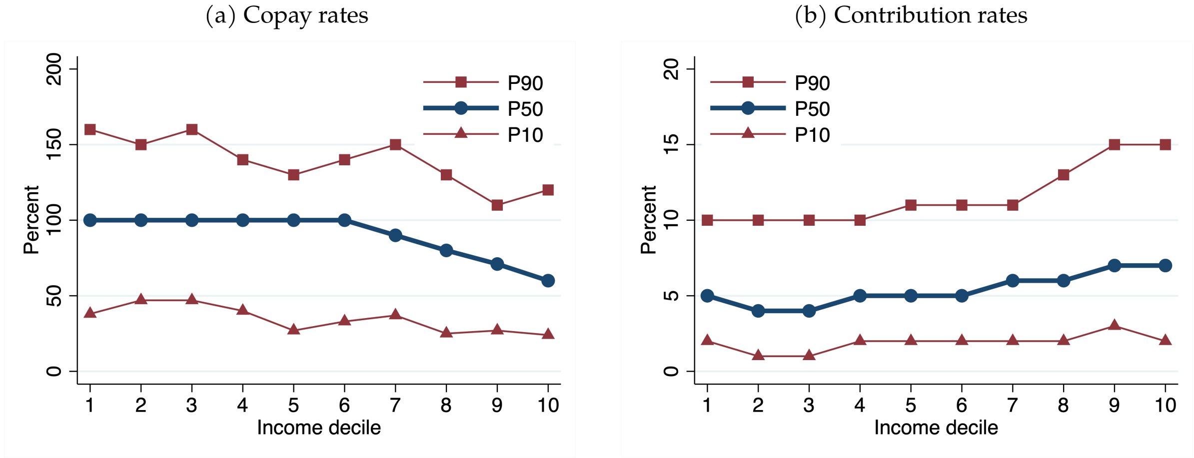

Notes: Panel (a) shows the distribution of copay rates by income deciles. The distribution is approximated by the decile-specific median, 10^{th}, and 90^{th} percentiles. Panel (b) shows the equivalent distribution for contribution rates. The results are based on the SCF 2010–2019 datasets. See Appendix A.1 for the sample selection criteria.

Figure 19(a) shows the distribution of copay rates by income. We approximate this distribution by reporting the median, the 10^{th}, and the 90^{th} percentiles of the income-group-specific distribution. Looking first at the median, we find copay rates to vary little with income for most households. For the bottom 60%, median copay rates are at 100% and they decline linearly towards the upper part of the distribution. Looking at the 10^{th} and 90^{th} percentiles, we find that the variation within income groups is substantially larger than the variation across income groups. We abstract from this heterogeneity in the model and assume a uniform copay rate. Such heterogeneity will be a driver of within-income-group inequality. Incorporating this additional heterogeneity would further increase wealth heterogeneity conditional on income. As we discuss in Section 5, this within-income-group inequality is currently too low in the model compared to the data.

Besides copay rates, the SCF data also provides information on contribution rates to employer-provided pension plans. Figure 19(b) shows the distribution of contribution rates by income. Again, we summarize this distribution by reporting the median, 10^{th} and 90^{th} percentiles.

Looking at the distribution by income groups, we observe, as for copay rates, that the within-income-group variation is much larger than the between-income-group variation, almost by an order of magnitude. However, contrary to copay rates, median contribution rates are slightly increasing with income, while we assume a uniform contribution rate by income in the model.

[33] Choukhmane et al. (2024) use a large employer–employee linked dataset and document large heterogeneity in contribution and copay rates by income and race.

Note that copay rates are declining with income, whereas contribution rates are increasing with income. If we consider the total inflow of savings by income, the combined copay and contribution flows, we find the two trends by income almost cancel each other as the contribution rate is almost 50% higher in the top decile, but the copay rate is 50% lower. Our model assumption of constant contribution and copay rates, therefore, matches the *total* inflow, which is important for the differences in saving rates. However, our model underestimates the *worker*'s share of the contributions at the top of the income distribution.

Similarly, one reason for heterogeneity in participation rates in home equity market, documented in Figure 2(a) in the main text, could be differences in mortgage interest rates. To study mortgage interest rates, we rely on HMDA data (see Section A.5.3 for details). In the HMDA data, we only observe interest rates for *approved* mortgages, but these data are still informative about how mortgage rates differ with income and if there is selection on mortgage rates. If approval rates depend on the income-interest rate relationship in a way so that low-income households only get mortgages approved if the interest rate is sufficiently high, then we should see a distribution of mortgage interest rates shifted towards higher values for lower-income households.

Figure 20: Mortgage interest rates in the HMDA data.

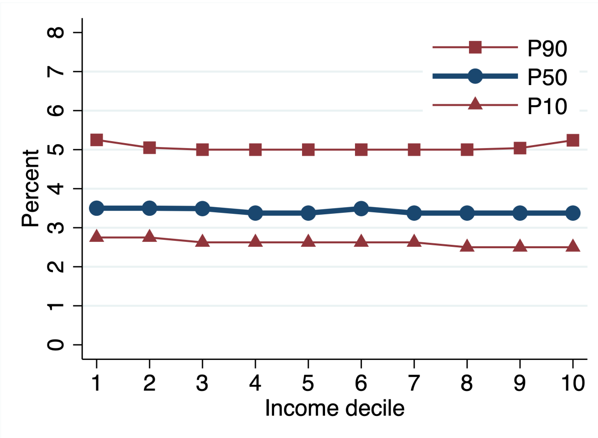

Notes: Distribution of interest rates by income deciles for approved mortgages. The distribution is approximated by the decile-specific median, the 10^{th}, and the 90^{th} percentiles. The results are based on the publicly available dataset spanning the survey years 2018, 2019, 2020, and 2021. The sample is restricted to applicants aged 25–64 with positive income.

Figure 20 shows the distribution of mortgage interest rates by income. To describe the distribution, we show the median mortgage rate and the 10^{th} and 90^{th} percentiles of the income-group-specific mortgage interest rate distribution. Two observations are important for our analysis. First, there is very little variation in interest rates at the median along the income distribution. The median interest rate is roughly at 3.5% along the entire income distribution. This finding supports our assumption of a uniform mortgage interest rate in the model. Second, the within-income-group dispersion is orders of magnitude larger than the between-income-group variation, but the support of the interest rate distribution varies little with income. The constant support of interest rates by income suggests that there is little selection on interest rates at the approval margin. However, the large dispersion conditional on income supports the idea that the within-income-group wealth dispersion is driven by return differences, which we abstract from in our analysis.

A.5.3 Home Mortgage Disclosure Act

To corroborate our findings regarding the mortgage approval rates by income obtained based on the SCF data (Section 2.3), we use the Home Mortgage Disclosure Act (HMDA) data, which provides a comprehensive data set on the US mortgage market. The HMDA data is published on an annual basis by the Consumer Financial Protection Bureau. Importantly, it includes the data on loan-to-value ratio and debt service-to-income ratio.

Sample selection

The available dataset spans the years 2018 to 2021, covering four survey waves. We restrict the sample to applicants aged 25–64 with positive income. To remove outliers, we exclude observations where the applicant is in the top 5% of the income distribution or the loan amount falls into the top 5% of the loan amount distribution. Additionally, observations with missing data for either sex or race are excluded. We classify a mortgage application as approved if the loan was originated or if the loan application was approved but not accepted. Conversely, we classify an application as rejected if the loan application or preapproval request was denied by the financial institution. Our sample includes loan applications for all purposes, but our results remain robust if we restrict the sample to mortgage applications or to both mortgage applications and mortgage refinancing.

Variable definitions

Some of the main variables exploited in our analysis are as follows:

- *Application Approved* – An indicator variable that equals one either if the loan was originated or if the loan application was approved but not accepted by the applicant. If, instead, the loan application was denied by the financial institution or the preapproval request was denied by the financial institution, the indicator variable equals zero. The indicator variable is constructed from the responses captured in the categorical variable *action_taken*.

- *Age Group* – A set of indicator variables that indicate the applicant's age. Age groups are 25–34, 35–44, 45–54, 55–64.

- *Debt-to-Income Ratio* – An indicator variable that equals one if the applicant's debt-to-income ratio is above 35% and zero otherwise. This indicator variable is constructed from the categorical variable *debt_to_income_ratio*, defined as the applicant's total monthly debt service to total monthly income.

- *Income Decile Dummies* – A set of indicator variables that indicate the applicant's income decile. The first income decile constitutes the reference group for the probit regressions.

- *Loan-to-Value (LTV) Ratio* – A continuous variable that captures the ratio of the total amount of debt secured by the property relative to the value of the property for which the loan is taken out.
- *State* – A set of indicator variables that indicate the applicant's state in which a given loan application was made.

Results

Table 11 presents the probit estimation results of mortgage approval rates, with the last column being our prefererd specification since it controls for loan-to-income ratios and debt-to-income ratios. We confirm that conditional on various household characteristics, income deciles exert a statistically significant and positive effect on mortgage approval rates, consistent with our findings from the SCF data.

A.5.4 National Compensation Survey

We corroborate the results regarding the heterogeneity of offer rates for employer-sponsored retirement plans by income, documented in Figure 4(a) in the main text, using data from the National Compensation Survey (NCS) on employer-provided pension plans conducted by the Bureau of Labor Statistics. This publicly accessible dataset includes offer and pick-up rates computed for selected percentiles of the wage distribution. For more granularity, we have acquired additional results categorized by income deciles from a special request to the BLS. Our analysis is confined to the period from 2010 to 2020, aligning with our examination of the SCF data. Additionally, we limit our focus to full-time workers employed in the private industry.[34]

Figure 21 presents the results. Two of the findings obtained in the SCF sample (Section 2.3) continue to hold: We observe a large variation in the offer and pick-up rates by income, with the variation in the offer rates playing a more dominant role.

[34] The series we use is "Percent of private industry workers with access to retirement benefits: all retirement plans; for average wage category is within X percent", where X refers to a given percentile, and "Take-up rate for private industry workers with access to and participating in benefit retirement benefits: all retirement plans; for average wage category is within X percent".

Table 11: Probit regression of approval rates for mortgages (HMDA).

	(1)	(2)	(3)	(4)
Income decile:				
-2nd	0.439	0.424	0.384	0.366
	(0.001)	(0.001)	(0.001)	(0.001)
-3rd	0.562	0.551	0.497	0.466
	(0.001)	(0.001)	(0.001)	(0.001)
-4th	0.631	0.622	0.559	0.515
	(0.001)	(0.001)	(0.001)	(0.001)
-5th	0.697	0.693	0.624	0.568
	(0.001)	(0.001)	(0.001)	(0.001)
-6th	0.736	0.741	0.666	0.596
	(0.001)	(0.001)	(0.001)	(0.001)
-7th	0.793	0.794	0.715	0.627
	(0.001)	(0.001)	(0.001)	(0.001)
-8th	0.836	0.846	0.763	0.658
	(0.001)	(0.001)	(0.001)	(0.001)
-9th	0.865	0.886	0.801	0.677
	(0.001)	(0.001)	(0.001)	(0.001)
-10th	0.872	0.904	0.819	0.668
	(0.001)	(0.001)	(0.001)	(0.001)
Loan-to-value ratio			-0.000	-0.000
			(0.000)	(0.000)
Debt-to-income ratio $> 35\%$				-0.310
				(0.001)
Demographic controls		x	x	x
State & Year controls		x	x	x
Observations	37,819,686	37,819,684	35,097,385	35,097,385

Notes: The regressions are based on the publicly available dataset spanning the survey years 2018, 2019, 2020, and 2021. The sample is restricted to applicants aged 25—64 with positive income. Standard errors are reported in brackets below each coefficient estimate. A loan application is approved if the loan was originated or if the loan application was approved but not accepted by the applicant. A loan application is rejected if the loan application is denied by the financial institution or the preapproval request is denied by the financial institution. Loan applications for all purposes are considered. All regressions include a constant. Demographic controls include age, gender, and race. See the sample selection criteria and the variable definitions in Appendix A.5.3.

Figure 21: Offer rates and pick-up rates for employer-provided retirement accounts (NCS).

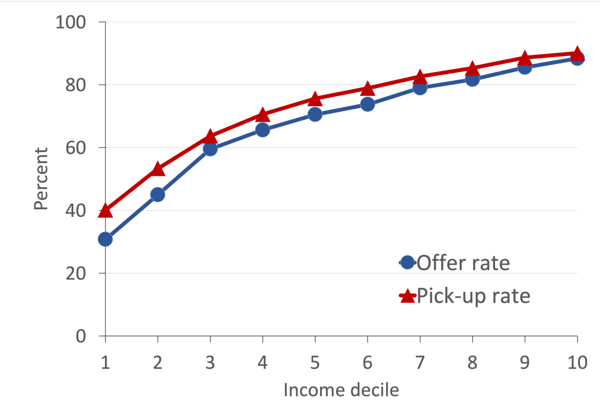

Notes: The figure is constructed using (partially) publicly available data from the National Compensation Survey (NCS) on employer-provided pension plans for 2010–2020. Tabulation by wage income deciles was provided upon a special request to the BLS. *Offer rate* denotes the fraction of employees whose employer offers a retirement plan and who are eligible to participate. *Pick-up rate* is defined as the proportion of employees who actually enroll in the employer-sponsored retirement plan out of the total pool of employees who have access to such a plan and are eligible to enroll.

B Calibration

B.1 Stochastic process for labor income and business ideas

Figure 22(a) shows that the model matches the empirical variance profile of log income for workers by age in the SCF. Figure 22(b) shows that the model accurately fits the age profile of business wealth-to-income ratios.

Figure 22: Calibration of the stochastic process for labor income and business ideas.

(a) Variance of log labor income by age. (b) Business wealth-to-labor income by age.

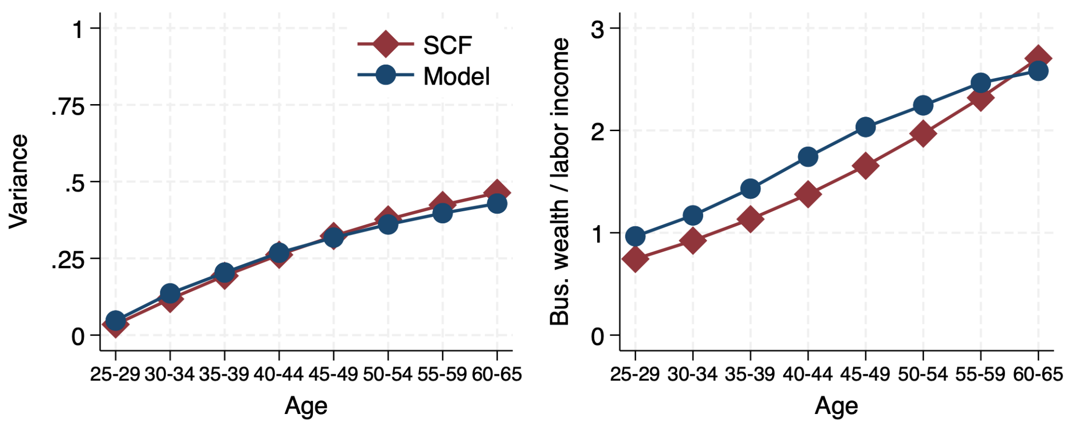

Notes: In Panel (a), the raw variance profile in the data is smoothed by fitting a quadratic polynomial in age. The intercept of the estimated variance profile in the data is set to zero. In Panel (b), we construct the profile in the model and data by dividing the average business wealth among business owners of a given age by the average labor income among workers of the same age.

B.2 Income-dependent financial frictions

Figure 23 displays the probability schedule of mortgage approval rates $\pi_H(y)$ (Panel (a)) and the schedule of (flow) offer rates for employer-sponsored retirement accounts $\pi_P(y)$ (Panel

Figure 23: Calibration of the financial frictions.

(a) Mortgage approval rates (b) Offer rates

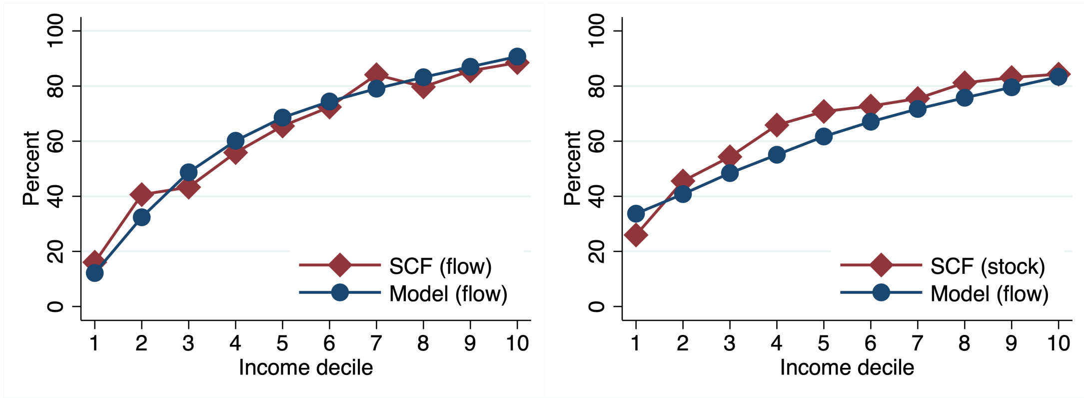

Notes: Panel (a) shows the mortgage approval rates in the SCF data and the model. The data line is taken from Figure 4(b) and represents the flow of approved mortgage applications in the past 12 months. The model line corresponds to the approval rates $\pi_H(y)$, which represent the flow of households whose mortgage application is approved in each period. We fit a functional form $\pi_H(y) = \alpha_0 + \alpha_1/y$ to the approval rates in the data. Panel (b) shows the offer rates for employer-sponsored retirement accounts in the data and the model. The data line is taken from Figure 4(a) and represents the stock of households having received an offer. The model line corresponds to the offer rates $\pi_P(y)$, which represent the flow of households receiving an offer in each period. We assume that $\pi_P(y) = \beta_0 + \beta_1/y$ and calibrate (β_0, β_1) inside the model to match the stock offer rates in the data.

(b)) used as inputs in our model. For mortgage approval rates, we fit outside of the model a functional form $\pi_H(y) = \alpha_0 + \alpha_1/y$ to the approval rates in the SCF data. We specify the probability to be offered an employer plan to follow $\pi_P(y) = \beta_0 + \beta_1/y$ and estimate parameters (β_0, β_1) inside the model to match the stock of offer rates observed in the SCF dataset.

To justify specifying the probability of being offered an employer plan as a function of the agent's income only, we present Figure 24(a), which shows the marginal effects of income deciles on offer rates obtained from the probit regression analysis of Appendix A.5.1. We report the marginal effects based on the full sample as well as three subsamples depending on the household head's age. The figure demonstrates that conditional on household income, there is no large variation in offer rates for employer-sponsored pension plans by age in the SCF, which is why we chose to specify offer rates in the model, $\pi_P(y)$ as a function of household income only.

Figure 24(b) provides similar evidence for mortgage approval rates using the SCF data.

Finally, since we observe the applicant's debt-to-income ratio in HMDA data, we construct Panel (c), which reports the marginal effects of income deciles on mortgage approval rates obtained from the probit regression analysis of Appendix A.5.3. We plot the marginal effects from a full sample as well as two subsamples based on the applicant's debt-to-income ratio, where we distinguish between high debt-to-income ratio (above 35%) and low debt-to-income ratio (below 35%). The figure demonstrates that, conditional on income, debt-to-income ratios do not substantially affect the probability of being approved for a mortgage. This serves again as a justification for why we chose the probability of being approved for a mortgage to depend on the agent's income only in the model.

Figure 24: Marginal effects of income on offer rates and mortgage approval rates in the SCF.

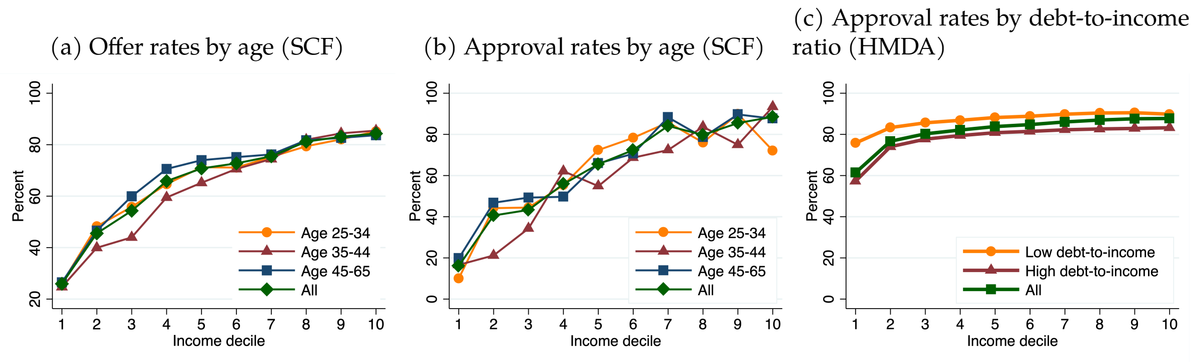

Notes: The figure illustrates the marginal effects of income decile dummies on two probabilities: being offered a retirement plan (Panel (a)) and being approved for a mortgage (Panels (b) and (c)). Panels (a) and (b) are derived from the SCF data detailed in Appendix A.5.1, with effects estimated from a full sample and three subsamples based on the household head's age. Panel (c), based on HMDA data (Appendix A.5.3), estimates effects from a full sample and two subsamples categorized by the applicant's debt-to-income ratio: *Low* (below 35%) and *High* (above 35%).

C Findings

C.1 Model fit

Figure 25 shows that the baseline model fits well the (unconditional) wealth-to-income ratios for each wealth component of core wealth in the SCF data.

Table 12 reports the shares of each core wealth component as well as the share of core wealth held by each income quintile in the model and data.

Table 12: Joint distribution of wealth and income in the model and data.

	Income quintile				
	First	Second	Third	Fourth	Fifth
	A. Core wealth				
SCF	1.84	4.84	8.80	14.99	69.52
Model	1.30	3.98	9.52	22.23	62.97
	B. Pension wealth				
SCF	0.82	3.51	8.10	16.84	70.73
Model	0.63	3.33	9.46	22.58	64.00
	C. Housing wealth				
SCF	2.79	6.44	10.45	16.13	64.18
Model	1.79	4.86	10.54	22.67	60.14
	D. Business wealth				
SCF	1.11	2.31	4.04	7.51	85.04
Model	1.62	2.48	5.12	18.97	71.82

Notes: The table reports the shares of wealth components held by each quintile of the income distribution in the model and SCF. The reported values take non-participants in the respective asset market into account.

Figure 26(a) contrasts the mortgage application rates by income in the SCF data and the model. Consistent with the data, the model endogenously generates a flat profile of application rates

Figure 25: Wealth-to-income ratios for wealth components in the model and SCF data.

Notes: The SCF sample covers periods 2010–2019. The value of each asset component in the data is computed net of debts. See Appendix A.1 for the sample selection criteria.

across the income distribution.[35] Figure 26(b) shows the pickup rates for employer-sponsored retirement accounts. These rates are defined as the fraction of workers with active pension accounts relative to workers without an active account but whose employers offer them the opportunity to participate. Following the strategy in the empirical section, we restrict the sample in the model to workers but keep the income deciles the same as in the full sample for the sake of comparability. The model successfully matches a rising enrollment rate with income, generating a monotone demand for retirement wealth.

Figure 27 shows the evolution of wealth-to-income ratios by age, which is an additional non-targeted dimension of wealth accumulation. The model achieves a good fit along this dimension. Models of life-cycle wealth accumulation typically struggle to fit the almost linear increase of the wealth-to-income ratio as they combine a deterministic life-cycle income profile and frictionless access to asset markets. Households know that income grows later in life and, therefore, start accumulating wealth only when income is above the average deterministic life-cycle trend. It is the combination of pre-specified financial contracts and frictional market access in our model that leads to a constant accumulation of savings starting early in the life cycle.

[35]See discussion to Figure 4 in the empirical section for details on the empirical measurement.

Figure 26: Demand for housing and retirement accounts in the model and data.

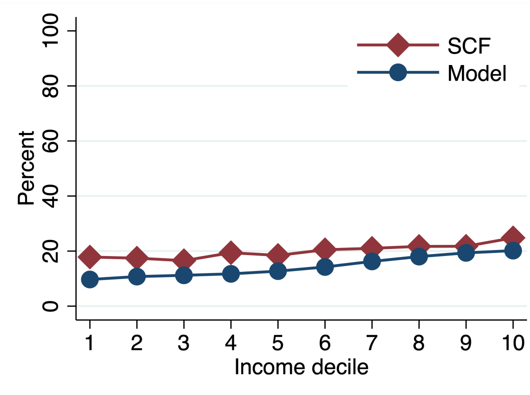

(a) Mortgage application rates

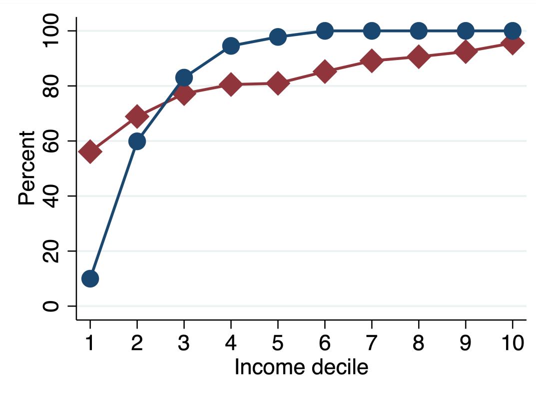

(b) Pick-up rates for retirement accounts

Notes: In Panel (a), an applicant in the SCF is a household who submitted a loan application in the past 12 months or who did not submit an application in the past 12 months because he/she thought it would be denied. Applications for a mortgage, home-based loan, and refinancing are considered. In Panel (b), *Pick-up rate* is defined as the proportion of employees who enroll in the employer-sponsored retirement plan out of the total pool of employees who have access to such a plan and are eligible to enroll.

Figure 27: Wealth-to-income ratios for wealth in the model and data.

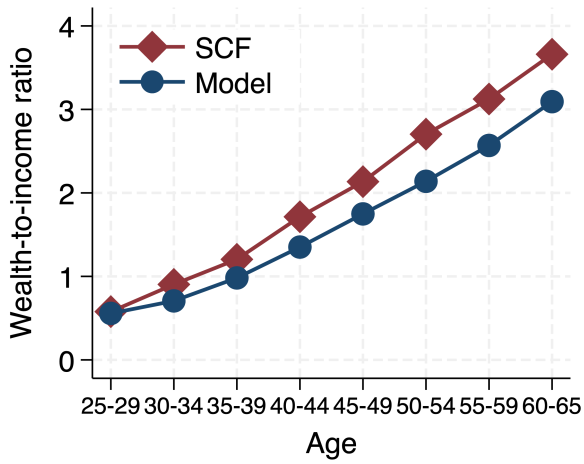

C.2 Frictionless access to the mortgage market

In the second experiment, we study the portfolio changes if we unilaterally improve market access in only one asset market. This experiment provides important insights to understand the results of the policy analysis below, where only the frictions in the pension market are changed. Here, we consider the mortgage market where we set the probability of being approved for a mortgage, $\pi_H(y)$, to one for all income levels so that all households who apply for a mortgage will also be approved. We keep access to the pension market and all other parameters at their calibrated values. In this case, it becomes easier for all households to accumulate wealth in the housing market, and hence, in relative terms, the frictions in the pension market increase. From a macroeconomic perspective, such changes could, for example, result from financial innovations in mortgage markets.

Figure 28 shows the percentage change in wealth-to-income ratios across income deciles in this

Figure 28: Change in wealth-to-income ratios in the experiment with frictionless access to the mortgage market.

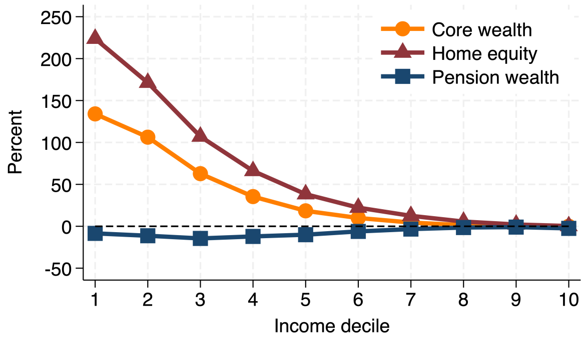

Notes: The figure shows the percentage change in wealth-to-income ratios for wealth, home equity, and pension wealth in the homogeneous access experiment relative to the respective ratio in the baseline model. The experiment sets the mortgage approval rates to one for all income levels. Wealth defined in eq. (2) is the sum of home equity, business wealth, and retirement accounts.

experiment. Comparing it to Figure 11a, where we removed only the *heterogeneity* of market access in the mortgage and the pension market, we now find asset substitution from pension wealth to home equity, especially in the bottom half of the income distribution, where access to asset markets is lowest. Below-median income households reduce their pension wealth-to-income ratio by 11% while simultaneously increasing their home equity substantially (121%). On the net, they increase their wealth by almost 70%. In the upper half of the distribution, the gross changes are insignificant, which shows that high-income households have sufficient saving opportunities. To summarize, wealth accumulation in the upper half of the distribution is not affected, but the shift in the asset market access leads to asset substitution towards the asset with lower frictions for income-poor households, thereby changing their composition of wealth.